


CONNEXIONS



The Interoperability Report

February 1996

Volume 10, No. 2

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

| | |
|-------------------------------|----|
| ISDN for Internet Access..... | 2 |
| Top Ten Web Trends..... | 9 |
| What does the IAB Do?..... | 14 |
| Residential Broadband nets.. | 18 |
| Announcements..... | 31 |

ConneXions is published monthly by Interop Company, a division of SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California, 94404-1138, USA.

Phone: +1 (415) 578-6900

Fax: +1 (415) 525-0194

E-mail: connexions@interop.com

Subscription hotline: 1-800-575-5717
or +1 610-892-1959

Copyright © 1996 by Interop Company.
Quotation with attribution encouraged.

ConneXions—The Interoperability Report
and the *ConneXions* logo are registered
trademarks of Interop Company.

ISSN 0894-5926

From the Editor

If you access the Internet from your home computer using a modem you have no doubt been dreaming about faster access methods such as ISDN or better. This month we bring you two articles describing the state of ISDN and the future use of cable TV for Internet access.

Integrated Services Digital Networks (ISDN) has been around for a number of years, but is still not fully deployed or available in many parts of the US. One perspective on this technology can be found in a recent issue of David Strom's *Web Informant*: "According to lots of trade press, ISDN is ready to take off as a speedy Internet access on-ramp solution. After trying it myself, I have to state flatly that it ain't gonna happen until phone companies and Internet providers work more closely together. Sure, it is fast: about five times the effective throughput of a 28.8 connection. But *ISDN Is Still Difficult to Network* properly, it still doesn't work well, standardization efforts are just getting started, and there are too few Internet Service Providers around that offer ISDN access. If you pass all these hurdles, it can be a very effective way to access the Internet... If you are thinking about doing this, you'll need either a lot of spare time, a great deal of patience and persistence, or a good consultant." Our first article by Dory Leifer is an exposition on some of the technical issues related to ISDN for Internet access. You can learn more about this topic by attending his ISDN tutorial at the next NetWorld+Interop in Las Vegas in early April. For more information or to register call 1-800-INTEROP today.

The *World-Wide Web* (WWW) is perhaps the single largest revolutionary development to hit the Internet. With easy-to-use browsers and thousands of "places to go," the Web has made the Internet accessible to the masses. David Strom, in addition to being a frustrated ISDN user, also writes about Web developments. We asked him to give us a perspective on trends and technologies in the Web.

ConneXions has published several articles about the *Internet Architecture Board* (IAB) in the past. However, we have not done so since 1989, and many aspects of the IAB have changed. This month we bring you an updated "insider's view" of the IAB. The article is by the current IAB Chair, Brian Carpenter.

The next logical step for Internet access to the home is to use the cable television system. Clearly such access would give users much higher bandwidth than either high speed modems or ISDN. In our final article, Mark Laubach gives an overview of developments in the area of residential broadband networking.

ISDN for Internet Access

by Dory Ethan Leifer, University of Michigan

Introduction


Five years ago I published an article for *ConneXions*, "ISDN: Why use it" describing some of our experiences at the University of Michigan and advocating the spread of Internet access via this new telecommunications service [1]. Since then, ISDN has been one of the most promoted as well as maligned service in public network history. However, during the last year, it is becoming clear that not only is ISDN here to stay but likely will change the landscape of inter-networking.

Integrated Services Digital Networks (ISDN) is now available in most U.S. metropolitan cities and in many developed areas across the globe. Current estimates put the global total at about two million lines with about a half a million in the U.S. While this represents a very small fraction of the 150 million U.S. telephone subscriber lines, ISDN line installations are nearly doubling on an annual basis. In the United States, ISDN is being fueled by demand for low-cost digital access to corporate networks as well as access to the Internet. Residential access is one of the most important growth areas for ISDN services used for telework, education and entertainment.

Are there better options other than ISDN to address residential access? Ultimately, the answer is yes. ISDN's lineage is clearly from the public switched telephone network (PSTN). ISDN 64Kbps channels are several orders of magnitude slower than modern LANs and the fact that most ISDN services are connection-oriented requires explicit call setup and tear-down to access connectionless services. On the other hand, ISDN also inherits other features from the existing public switched network. The circuits are reliable once properly installed and can be installed in large number without some of the usual scaling issues that plague shared LANs. By far, the most important driver for ISDN now is the desire for low-cost digital access and ISDN is the only wide-scale available service that fits the bill.

Modems

Why not use modems? Modem speeds have increased dramatically during the last ten years to rates once never thought possible; why not wait for faster modems? Experts agree now that 28.8Kbps modems are approaching the theoretical limit of an analog or POTS (plain old telephone service) line. Real statistics from operating modem pools indicate that even 28.8Kbps is a stretch. Figure 1 illustrates the distribution of final negotiated modem speeds from 28.8 Kbps V.34 attempts; only a small fraction of the connections obtained the full 28.8Kbps.



Source: Merit Network, Inc. Modem Dial-in Service. Ann Arbor, MI

Figure 1: Sample Modem Connection Profile

ISDN is delivered in two flavors, basic and primary rate. The *Basic Rate Interface* (BRI) is based on a special ISDN-specific physical-layer protocol which is capable of carrying 2 circuit switched B-channels and a signaling D-channel about 5.5Km over a single pair of normal voice grade wire. The *Primary Rate Interface* (PRI) is designed to be carried over standard T1 or E1 circuits. PRI consists of a single D-channel and either 23 (T1) or 30 B channels (E1). Optionally, a D-channel can be shared across several PRIs to allow for an extra B channel on the T1 or E1.

Pricing for PRI and BRI varies dramatically between regions and countries. In some parts of Europe, a PRI can be ordered for about the same price as three or four BRIs while in regions of the United States, PRI is always significantly more expensive per B-channel.

How ISDN is used for Internetworking

For a variety of tariff-related considerations, ISDN lines are used in the United States principally as an access service and as either an access service or bona-fide WAN in other parts of the world. As an access service, ISDN is configured similarly to a dial-in modem service. At the edge of an existing network, a *Network Access Server* (NAS) is equipped with a number of available B-channels. These channels can be delivered by a combination of BRI and PRI depending on the number of users to be served and the local pricing of PRI vs. BRI. In almost all cases, ISDN dial-in services are used to deliver peer-to-peer or network nodal service, as opposed to dumb terminal access. The NAS supports the *Point-to-Point Protocol* (PPP) and is usually capable of routing IP or IPX in addition to transparent bridging.

One of the most flexible options on the NAS is support for what is commonly called "digital modems." It is possible for ISDN lines to interoperate with plain old telephone service. A call that initiates on a voice line can terminate on an ISDN line. Fundamentally this type of interoperability is required to make the service useful for voice. When a NAS receives a call on an ISDN line, it can either process it as a data call and perform HDLC framing on the data or it can process it as a modem call with modem protocols such as V.32 or V.34. The latter requires either modem chips in the NAS or powerful digital signal processors (DSPs). A NAS with modem support is capable of simultaneously functioning as an access server for digital dial-in as well as a replacement for traditional modem dial-in servers. Not only does this arrangement combine both functions into a single box, but it has the added benefit of delivering modem service on less error-prone ISDN digital circuits.

Remote ISDN-equipped locations usually use BRI because of cost sensitivity. A residence, for example may install BRI to attach some combination of computers, telephones and fax machines. Until recently, it was very difficult to economically justify the replacement of residential POTS for ISDN but now the case is much easier because of the availability of ISDN access equipment with built-in voice features as well as new telephone company tariffs that provide residential services on ISDN. Some ISDN users have even canceled their POTS service and have ISDN as their only telephone and data service. Issues with this arrangement are discussed shortly.

Ethernet

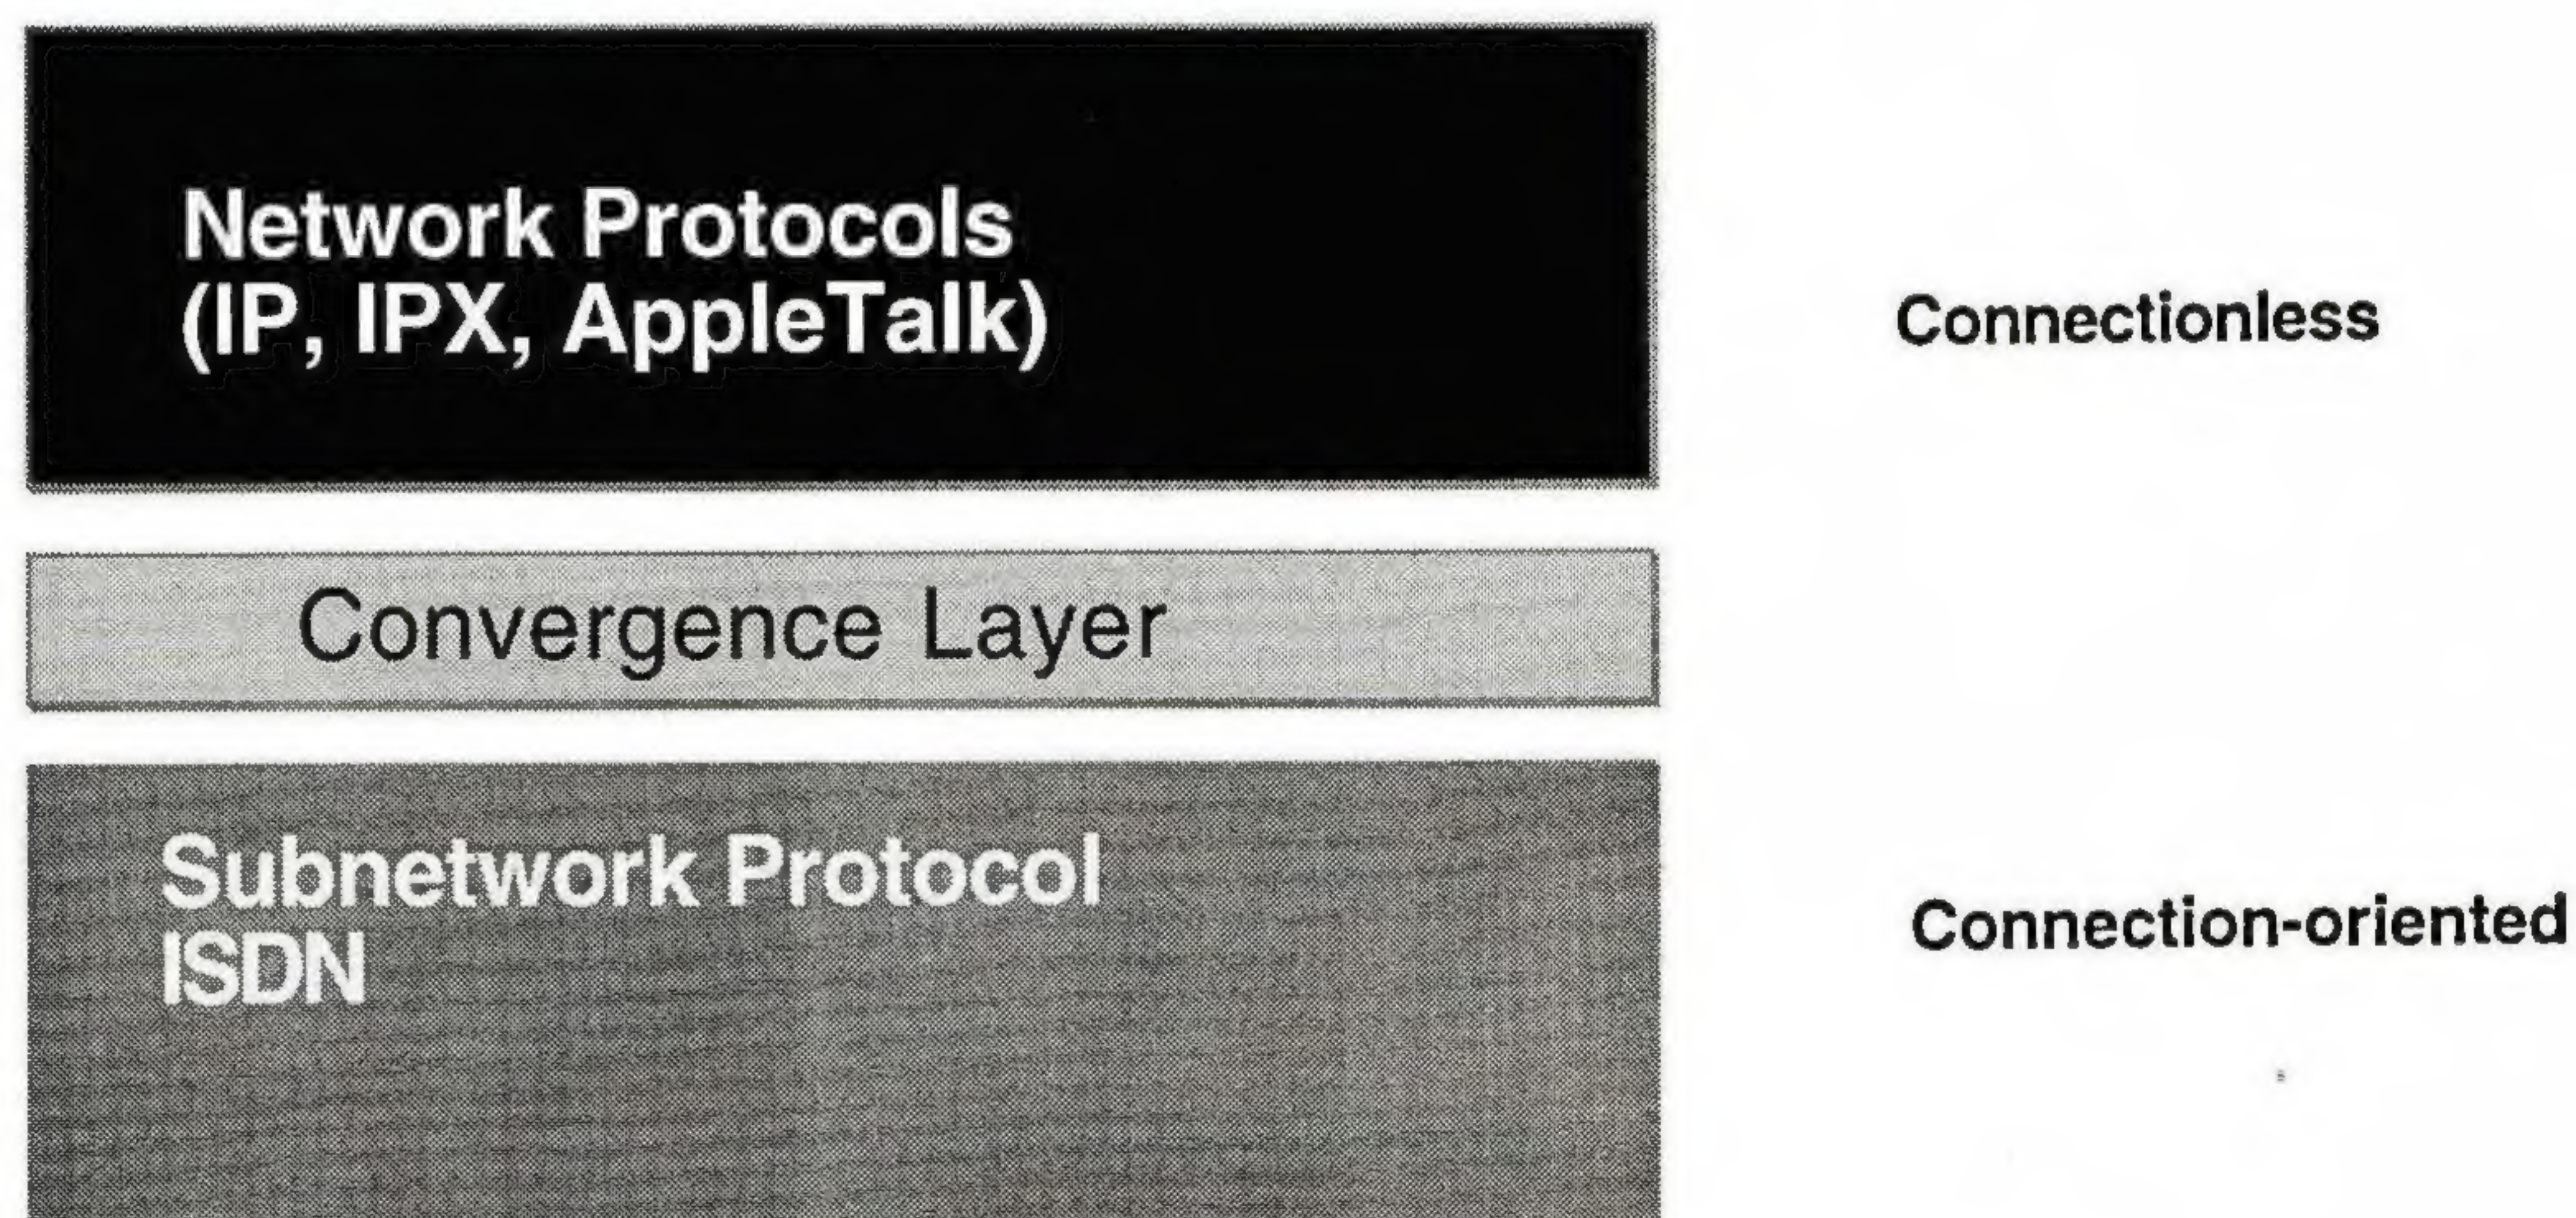
There are at least three distinct types of devices which can deliver bridged or routed networking to remote locations via ISDN. The first type is a device with Ethernet and usually a single BRI, commonly called an *ISDN Remote Bridge/Router*.

ISDN for Internet Access (*continued*)

This device is a pared-down version of larger enterprise routers with added smarts to handle the ISDN interface as well as connection management, discussed shortly. Since the interface to hosts is an Ethernet LAN, there is no need for special purpose software on the hosts.

TA The second type of device is a sync-to-async *Terminal Adapter* (TA). These devices attach to a host's serial port at up to 115.2Kbps. When used with PPP software on the host, this TA is able to convert async PPP on a single host to sync PPP on the ISDN B-channel. The TA is often referred to as an "ISDN modem" because it looks like a Hayes-compatible modem to the host software. In reality the term "modem" here can be misleading because the TA may not have any ability to handle modem protocols like V.32 and V.34 on the ISDN line. There are a number of new async-to-sync terminal adapters that are becoming price-competitive with POTS modems.

Internal Adapter The third device for remote ISDN access is the *Internal ISDN Adapter*. Available for most bus architectures, ISDN internal adapters are usually offered with software which implements either an ISDN-specific API such as WinISDN or the European CAPI, or a network interface such as NDIS or ODI. While the latter is designed to fit ISDN into the host LAN stack, the former is capable of supporting multimedia applications such as a voice response system. Some implementations of PPP are built on top of an ISDN API. [2]



LAN network protocols are connectionless. ISDN is connection-oriented. The Convergence Layer is always making best guesses.

Figure 2: ISDN Convergence Layer

Issues Connection management is one of the biggest technical changes for either routing or bridging over ISDN. Since ISDN is connection-oriented and IP and IPX are connectionless, a convergence function is required to handle the call setup and tear down. Paradoxically, this function needs to both maintain transparency so hosts and users "think" they are always connected to the network as well as minimize connection times to reduce toll charges and quickly free idle ports on the NAS. The term "bandwidth on demand" describes this capability. The trick is to leave the ISDN circuit disconnected until packets need to be exchanged across it; after the packets pass, disconnect the circuit. Consider, however, that several LAN network protocols, routing protocols and even applications exchange messages even with no user activity.

These messages will cause the ISDN circuit to either be brought up and down to accommodate the messages, perhaps every 30 seconds, or possibly worse, leave the call connected forever. There are a number of clever techniques to alleviate some of these woes; one example is *spoofing*.

Spoofing is used by ISDN routers and bridges to fool other routers and hosts to believe that the ISDN circuit is connected all the time. Each router or bridge synchronizes a database containing routes, services, active users, etc. with its peer on the other side of the ISDN call. The two parties then disconnect the circuit between them and start to proxy for each other's networks causing (hopefully) the networks to behave as if the circuit was still connected. Once in a while they call each other and exchange updates or else network state changes would never be reported. If everything works as planned, nobody will notice that a switched circuit is involved, and the network administrator will receive a modest telephone bill. In practice, it may be quite difficult to reliably eliminate all non-critical traffic. Consider what happens if someone leaves *xclock* open on their X-server!

Equipment interoperability has been a major issue for ISDN. Until about 1992 there were sufficient ISDN protocol variations to make one wonder whether a piece of equipment was going to be able to operate with a particular ISDN switch. With *National ISDN* in North America this is much less of a concern that the equipment can be made to *eventually* work [3]. Now telephone companies and equipment and service providers are trying to make sure that it works the first time. Of course, bridges and routers need to be compatible not only with the telephone network but with each other. To address-end-to-end compatibility, the Point-to-Point Protocol was chosen as a standard for internetworking on circuit-switched ISDN. [4]

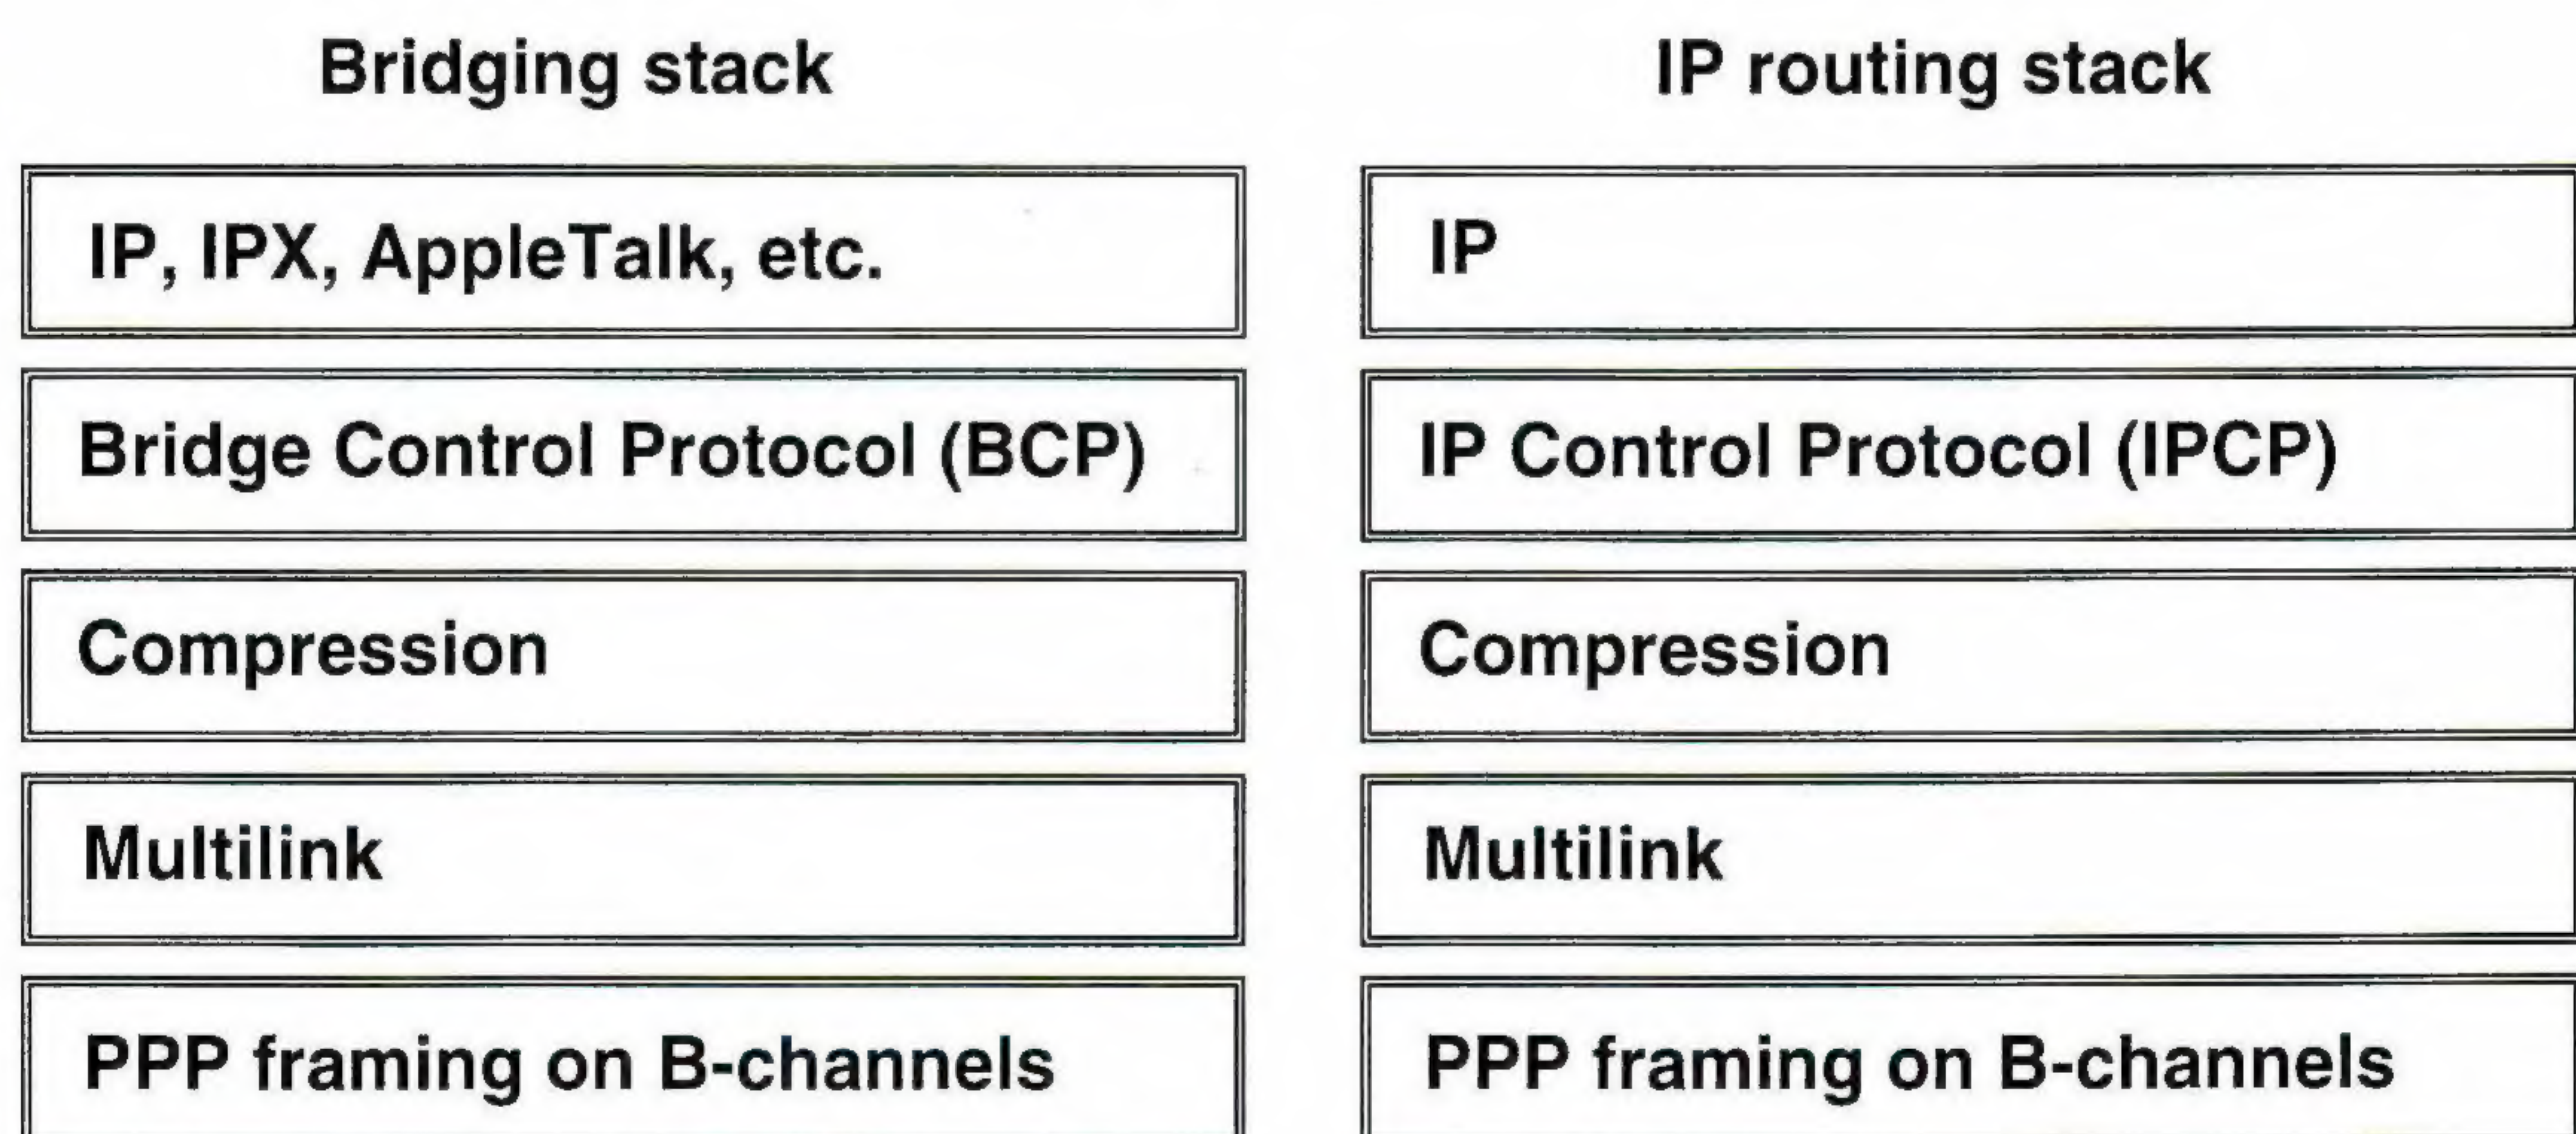


Figure 3: Typical PPP over ISDN Protocols

The Point-to-Point Protocol (PPP), [13] developed by the IETF has been implemented for most ISDN bridging and routing applications. PPP over ISDN is usually used on the B-channel in its sync HDLC form with CHAP (*Challenge Handshake Authentication Protocol*) for security. The *IP Control Protocol* (IPCP), *IPX Control Protocol* for NetWare (IPXCP), or *Bridge Control Protocol* for bridging functions (BCP) provides network functions. *Multilink PPP* (MP), one of the recent extensions to PPP, allows multiple B-channels to be joined together to make use of the full 128Kbps of a BRI line or aggregate higher bandwidths with multiple BRI or PRI channels.

continued on next page

ISDN for Internet Access (*continued*)

The technique involves starting PPP sessions on each B-channel and creating bundles of channels over which a single virtual PPP connection exists [5]. This type of channel aggregation is often and erroneously referred to as *BONDING*. *BONDING*, which stands for *Bandwidth ON Demand INteroperability Group* is a byte-synchronizing aggregation technique usually used for video conferencing although it can be used for routing or bridging.

For a variety of reasons, few of them technical, V.42bis data compression used for modems could not be as easily standardized for use with ISDN. As a result, there are a number of compression protocols to choose from. Peers need to negotiate the use of compression provided that they are able to support a common subset. There are several devices, for example, that use proprietary STAC compression in addition to or as an alternative to V.42

How to get an ISDN Internet connection

Step 1: Find out if you can order ISDN in your area. Unlike POTS that's available everywhere, ISDN is not. Although you may have seen numbers like 90% of subscribers can order ISDN, I seem to have a statistically disproportionate number of colleagues that are part of the 10%. Whether or not you can get ISDN is governed by at least two factors:

- How far do you live from the telephone company "serving office." If this distance is over 5.5Km (of cable, not air) and you are served with a copper loop, you will most likely require both repeaters as well as removal of loading coils designed to extend POTS. In newer developments, telephone companies install subscriber loop carrier systems (SLCs). SLCs are fed with fiber from a hosting central office and are placed close to subscriber concentrations such as in housing subdivisions. This effectively reduces the length of the copper loop and allows ISDN to be installed at virtually unlimited distances from the Central Office (CO). For loops that are completely copper and are longer than 5.5Km, mid-span repeaters, which regenerate the signal on route, are a possibility. Some telcos charge more for ISDN service if mid-span repeaters are required and cannot always notify customers at order time if additional charges will apply.
- Are you serviced by an ISDN-capable switch or is there one available in your "wire center?" Under some circumstances, telcos will delay installing ISDN upgrades to digital switches until demand warrants so. Just because you are connected to an AT&T 5ESS or Nortel DMS-100 (the two most popular digital switches in the U.S.) doesn't mean that you can get ISDN service. It may be necessary to foreign-exchange an ISDN connection from a near-by switch. There are often additional charges if this is necessary and your telephone number will be from the remote location. Telcos often refer to this arrangement as "virtual ISDN."

Step 2: Decide whether you want to use ISDN for both voice and data. Should you get rid of your home phone service when ordering ISDN? Generally, no. Although the "I" in ISDN is supposed to mean "integrated," many users are finding it more attractive to keep their POTS lines. A few reasons:

If the serving ISDN switch is the same switch that serves the POTS line, it is possible to move the number to the ISDN line without disrupting service. In practice, this is tricky business as ISDN installs are often not successful on the first try.

If the two mentioned switches are different, a number change is required which may not be very desirable. ISDN in the U.S. has no network power to drive ISDN equipment. This means that all power needs to come from the home or office. Prepare to have proper backup batteries to take over if you have only ISDN.

ISDN voice equipment is usually fairly rudimentary. Some ISDN equipment vendors have only recently started shipping voice interfaces. Take a long look at the voice features supported such as call-waiting, forwarding, hold, etc. and the quality of the telephone interface (adequate ring voltage) before deciding to use ISDN for voice. This situation will, of course, improve with product maturity.

Step 3: Choose a subscription type from an *Internet Service Provider* (ISP). Internet service over ISDN generally comes in one of three categories.

- *Dial-up, single host:* The user connects to the ISP over a PPP connection. The ISP usually dynamically assigns an IP address every time the subscriber connects which makes this service unsuitable for subscribers that wish to offer services like WWW servers. The configuration is exactly the same as dial-up analog service but provides more bandwidth and faster call setup times. The subscriber uses either an internal adapter or a device that looks like a Hayes-compatible modem, interprets asynchronous PPP on the DTE side and speaks synchronous PPP on the ISDN B-channel. Expect to pay ISP rates similar to POTS dial-in.
- *Dedicated network access:* In some cases, it is cheaper to use ISDN for full-time dedicated connections than standard leased line service. The term "dedicated" is a bit misleading because the ISDN service is still completely switched; calls are just rarely disconnected. This service is often used with Centrex "intercom" calls which have no per-minute or per-call telco charges. The ISP dedicates one or more B-channels for each of the subscribers so there is never a chance of contention or being blocked. Subscribers get static IP addresses and can run standard routing protocols if desired. Subscribers generally use an ISDN-equipped router or bridge to connect to the Internet.
- *Demand network access:* This mode of access is used on small LANs with part-time requirements for Internet access. When a host on the LAN attempts to send packets to a route pointing to the ISP, the router dials an ISDN call to the ISP and carries the traffic usually until an idle timer expires. This service is not well-suited for providing Internet services as ISPs may not be willing to initiate outbound calls when packets arrive from the Internet destined for the subscriber LAN. The ISP offering this type of service may decide to engineer it either for blocking or non-blocking. Addresses are assigned as static although there is interest in supporting dynamic addresses. There are a number of issues which need to be solved before remote LAN dynamic addresses can be supported. This capability may eventually be supported by the *Dynamic Host Configuration Protocol* (DHCP). [11, 14]

Conclusions

After a near disappearance in the press, ISDN has finally emerged as a popular service to provide inexpensive access to the Internet and private networks. With the current rate of growth, ISDN may be installed in more than a quarter of American homes with PCs within the next five years.

ISDN for Internet Access (*continued*)

There are still a number of important service as well as protocol issues that need to be understood by users and service providers before the service will scale comfortably to these levels, but the ISDN train has certainly left the station and is moving along quickly.

References

- [1] Leifer, D., "ISDN: Why use it?," *ConneXions*, Volume 4, No. 10, October 1990.
- [2] Ben-Artzi, Fedor, Klingman, Koren, Giessler, "WinISDN Specification," February 1995.
- [3] Bellcore Special Report NWT-002006, "National ISDN", Issue 1, August 1991.
- [4] Simpson, W., "PPP over ISDN," RFC 1618, May 1994.
- [5] Sklower, K., Lloyd, B., McGregor, G., & Carr, D., "The PPP Multilink Protocol (MP)," RFC 1717, November 1994.
- [6] Blackshaw, R., "Components of OSI: Integrated Services Digital Networks (ISDN)," *ConneXions*, Volume 3, No. 4, April 1989.
- [7] Hobbs, J., "Enterprise Networking and National ISDN Synergy," *ConneXions*, Volume 7, No. 12, December 1993.
- [8] Minoli, D., *Enterprise Networking: Fractional T1 to SONET, Frame Relay to BISDN*, Artech House, 1992.
- [9] "National ISDN-1," Bellcore Special Report NWT-001937, February 1991.
- [10] "National ISDN-2," Bellcore Special Report NWT-002120, May 1992.
- [11] Allard, J., "DHCP—TCP/IP Network Configuration Made Easy," *ConneXions*, Volume 7, No. 8, August 1993.
- [12] Lippis, N., & Herman, J., "Widening Your Internet Horizons: Wide-Area Options for Internets," *ConneXions*, Volume 5, No. 10, October 1991.
- [13] Hobby, R., "The Point-to-Point Protocol (PPP). A new proposed standard Serial Line Protocol," *ConneXions*, Volume 4, No. 4, April 1991.
- [14] Droms, R., "Dynamic Host Configuration Protocol," RFC 1541, October 1993.

DORY ETHAN LEIFER pioneered the first implementation of Internet service over ISDN at the University of Michigan in 1989. Since then, he has consulting for end users, telephone companies and equipment vendors interested in data networking over ISDN. He teaches the NetWorld+Interop "Internetworking ISDN" tutorial as well as seminars at Michigan on TCP/IP, ISDN and ATM technologies. Mr. Leifer is currently involved in strategic planning for campus and regional networks and holds an appointment with the Merit Network. He is involved with several wide-scale ISDN Internet access projects. He holds an M.S.E in Industrial and Operations Engineering from the University of Michigan and a B.S. in Computer Science from Rensselaer Polytechnic Institute. E-mail: leifer@terminator.rs.itd.umich.edu

Top Ten Interesting Trends on and around the World-Wide Web

by David Strom

Introduction

Writing about the Web for a print publication is quite a challenge: things change so fast that sometimes one's research is out of date not only before publication but before one's article is even saved on the word processor.

Our esteemed editor asked me to do an article looking at home pages in the style of Siskel and Ebert (I paraphrase his request). I'm not so sure what this means, but I will attempt to give you a brief look at where things are going in terms of Web developments from my perspective as well as some of the things I'd like to see happen with the Web over the next few years.

Trends

First, my top ten list of Web trends:

1. *Caffeine and anti-caffeine:* For purposes of discussion, let's call all efforts to combine the Web with multimedia, animation, agents and complex images under the heading of caffeine, and other efforts to combine the Web with proper search engines, directories of links, and mostly text-based methods under the heading of anti-caffeine. Both "sides" if you will are important for the Web to grow and thrive. But developers for the most part have chosen one direction or another for Web-based products. Most of the trade press has focused their attention on the caffeinated side of things, which is a Bad Thing: I believe the more interesting developments will be on the anti-caffeine side. But perhaps I am biased: I also prefer decaffeinated beverages.

Nevertheless, these caffeinated applications are important, if nothing because they free us from depending on a particular computing platform and also allow for more network-centric computing and program development. Now, do I believe this will happen anytime soon? Nope.

2. *Netscape and anti-Netscape:* The general public loves a two-sided battle, and one of the problems was that until recently the Web had so many sides that it was difficult to keep track. Now, thanks to Netscape, it is much easier: we have Netscape versus Everyone Else, or Netscape versus Microsoft. (You can fill in your favorite vendor if you'd rather.)

Netscape has had some positive effects on the Web: it has brought about a rapid deployment of graphical (versus text-based) browsers, innovation with respect to HTML tags and servers, and improved graphical look and feel of Web pages themselves ("this page has been optimized for ..."). It has brought about some negatives as well: a rapid deployment of overwrought graphical content, too frequent updates of browser software, and too many new tags and server extensions. All of this contributes to what I feel is the end of openness and a standards-based Web [1], which is a Bad Thing.

3. *Five million channels and nothing on:* First everyone on the Internet had their own (potential) home page. Then came along CompuServe, AOL and Prodigy and now those several million customers have their home pages too. Soon everyone will have their own Web server. Then what? Actually, this is a Good Thing: I believe the best role for the Web is to have a Web server into each desktop operating system, and by proliferating home pages (whatever they really are these days is hard to tell, but that's for another article) only helps to get this notion going.

continued on next page

Top Ten Interesting Trends on the Web (*continued*)

Why is this desirable? The kind of things that a Web server does (linking and organizing documents) is exactly the kind of things that a modern operating system should do, but doesn't—yet.

An interesting corollary to this trend is that everyone becomes their own publisher. This is not necessarily a Bad Thing. Take my own case as a shameless example. I publish my own e-mail newsletter now to a small audience of computer vendor marketing and engineering types. While I could do this before the rise of the Web, having a Web site to archive my back issues and provide links to other interesting places complements the overall publishing scheme nicely. The trick to being your own quality publisher is also finding the right mix of skills (editorial, production, sales and circulation) that are needed for the world of paper publishing as well.

4. *NT vs. UNIX*: The Web has been a great example demonstrating the best and worst of UNIX: if you know UNIX, setting up a Web server is no great stretch. The tools to administer a Web server have great power but obscure commands, and to the UNIX-ignorant are even more intimidating than the basic operating system. All fodder for the Windows NT camp. The Web and NT make beautiful music together: you can have your graphical user interface cake (something that all Windows users have taken for granted but that UNIX users only recently discovered) and eat reliable, rock-solid operations too (something that all UNIX users have taken for granted but all Windows users have only recently discovered).

NT Web servers are also for those corporations that have relative control over their desktops and are comfortable with (pre-NT versions of) Windows. There are now over a dozen NT Web servers [2], and I'm sure more are on their way. But more importantly than sheer numbers is the perception that NT makes a great Web server by the corporate buying public. My recommendation: if you don't know UNIX, don't start now for all things Web. Try one of the more popular NT Web servers.

But don't think that an NT box can completely replace all the functionality of a good UNIX server for all things non-Web (mail, network management, and netnews come to mind as three services that UNIX does particular well and NT is still far short on). In addition, some signs of immaturity still remain: NT Web server remote administration tools are bare-bones at best. Too few Internet Service Providers offer NT Web hosting services, although that will change quickly over time. And sometimes it can be confounding to have too many choices for software.

5. *New class of useless software—HTML editors and HTML add-ons to existing word processors*: The Web has brought about the fast demise of proprietary desktop word processing formats in favor of tagged text. Is this progress? A brief history of word processing, for those that might have forgotten: Back in the old (pre-PC) days, we had VT-100 terminals and control characters. TeX [3] if not king, was certainly a notable achievement in this venue. Then came Wang, NBI, Vydec, Xerox and other dedicated machines that did the job: you all know what happened to them. Out of Wang's popularity was born Multimate, which ran on PC DOS. They got acquired by Ashton Tate which was acquired by Borland, and then disappeared. Now we have three major vendors of word processors on the desktop.

Anyway, most of these HTML editors and HTML add-ons (Microsoft Word Internet Assistant, Lotus Word Pro, and WordPerfect's Internet Publisher) don't really add much to the process of creating, checking, and publishing HTML documents. My recommendation: find yourself a good text editor and return to the glory days of the past.

6. Latest buzz word—the "Intranet": We now have Attachmate saying they are "the Intranet company," a trade magazine with an entire "Intranet" section [4], and I'm sure trade shows with Intranet in their titles aren't far behind. In one week last fall, Microsoft, Sun, Netscape, and others made product announcements with the word figuring prominently. I've even written a white paper on the subject [5]. So what is Intranet, anyway? My take is simple: take one part TCP/IP, one part Internet technologies such as mail and Web, and one part corporate database, shake into a publishing metaphor and stir lightly into an HTML browser.

7. WEBng—database connectivity: Despite my attempts at humor, Intranets will become serious business as corporations figure out that client-server applications really mean access to data. The Internet technologies were nicely designed for this purpose, and expect to see more corporations co-opting this mix as they get further along in publishing more than just employee handbooks and corporate policy manuals on their internal Web servers. The ability to tie your mainframe and network databases and Web servers together will make this a viable business in years to come. We are just beginning to see products that make this possible, and again this is another Good Thing.

8. HTML as the new application interface: We now have a variety of "faceless" applications—software that has no interface of its own, but rather uses a Web browser instead to do user input and display information. Want to manage your network laser printers? [6] Use a Web browser. Want to check your calendar? [7] Use a Web browser. Want to manage your router? [8] Use a browser. I believe this is a Good Thing, but not because I want to run my Web browser for all my applications. Rather, this is desirable because it furthers the adoption of Internet technologies within the general-purpose office productivity market, and also makes it easier for me as a mobile office worker to move around the world and do my work over the Internet.

What will really help this along is TCP/IP. TCP/IP is now available as part of every new desktop operating system sold today. That's a big change—and a big improvement—from several years ago, when the protocol was only the province of the truly enlightened. But we still have several potholes along the IP highway to fill in: management tools, for example, that were designed for routers and not desktops; mixing earlier versions of Windows and IP is still somewhat of a challenge; and making Winsock and Open Transport work is full employment for an entire army of consultants.

9. Two words; Internet commerce: Depending on whom you talk to, either the Internet will never be safe for conducting commerce or it already is far and away safer than handing your happy waiter your credit card. Dick Shaffer of Technologic Partners sums it up best by saying "A standard payment mechanism for Net purchases will appear in 1996 (it will be called a 'credit card'). In the meantime, the various Internet *.cash systems will continue to duke it out."

Top Ten Interesting Trends on the Web (*continued*)

I got real insight into some of this process when I tried to buy some stuff via the Internet: shopping malls were confusing (not to mention relatively empty of patrons), incompatible forms with my browser, difficult to find stuff, overpriced shipping charges, etc. [9] All in all, a relatively unsatisfactory series of experiences. But this will get better, or else we will have to find some other use for the Web fast.

10. Pipes vs. content—the latest tectonic shift for ISPs: Something I've predicted a long time is that the traditional on-line access vendors (CompuServe et al) are moving to becoming providers of pipes rather than content. [10]. I don't mean to minimize the value of content on the on-line vendors at all, it is just that the Internet will receive more innovation and more interest. Already, AT&T and Apple have learned this lesson with Interchange and eWorld, respectively. This trend will continue, and we might even see phone companies figure out the former (but doubtful of the latter).

Wish list So, what would I like to see happen on, and around the Web?

- I'd like to see Microsoft stop talking about how good it is going to be and really try to integrate its Web services into NT.
- I'd like to see true 32-bit applications appear on Windows 95 and NT so that I wouldn't have to continually be reminded of the 8.3 character file names inherited from a 20-year old operating system.
- I'd like to see authoring and publishing tools work well with a wide range of Web servers so I wouldn't have to manually arrange my content and use an FTP client as my main organizational tool.
- I'd like to see a meaningful and stable HTML standard that was fully embraced by Netscape and Microsoft, so we can innovate in more useful ways than tag fights.
- I'd like to see the Macintosh reborn as a Web authoring platform, with all the necessary tools and integration, and ease of use that it had as a desktop publishing platform.
- I'd like to see the trade press focus on something other than caffeinated applets and agents, and examine the lack of meaningful ways to transport existing content onto a Web server.

And to put this all in the proper perspective, I'd like to see World Peace in my lifetime.

References

- [1] A more complete essay by the author on the end of openness, "Anybody Remember Open Systems On The Net?" can be found at: <http://techweb.cmp.com/techweb/docs/open.htm>
- [2] The World-Wide Web Consortium has a nice listing of these servers at <http://www.w3.org/hypertext/WWW/Servers.html>. Also, Paul Hoffman maintains a list of Web servers (along with a description of features) at: <http://www.proper.com/www/servers-chart.html>
- [3] *TeX and METAFONT*, Donald E. Knuth, Digital Press/American Mathematical Society, 1979. For places around the Web that offer the TeX product, check: http://www.yahoo.com/Computers_and_Internet/Desktop_Publishing/TeX/

- [4] *Interactive Week* magazine not only has an Intranet section, but an Intranet columnist: <http://www.zdnet.com/~intweek/>
- [5] "Creating Private Intranets: Challenges and Prospects for IS" by the author is available at:
<http://www.strom.com/pubwork/intranet.html>
- [6] Tektronix PhaserShare printer management software can be seen at http://www.tek.com/Color_Printers/support/demo550/
- [7] Campbell's OnTime calendaring software can be seen at:
<http://www.ontime.com>
- [8] Tribe's WebManage routers can be seen at:
<http://www.tribe.com/products/webmanage/tour.html>
- [9] An article by the author on the on-line shopping experience can be found in *Web Review* magazine at:
<http://gnn.com/wr/dec8/biz/olc/index.html>
- [10] See "In the On-Line Market, The Name of the Game is Internet," *New York Times*, 9/25/94.
- [11] Crowcroft, Jon and Handley, Mark "The World-Wide Web: How Servers Work," *ConneXions*, Volume 9, No. 2, February 1995.
- [12] Koster, Martijn, "Robots in the Web: threat or treat?" *ConneXions*, Volume 9, No. 4, April, 1995.
- [13] Berners-Lee, T., "A Summary of the WorldWideWeb System," *ConneXions*, Volume 6, No. 7, July 1992.
- [14] Berners-Lee, T., R. Cailliau, A. Loutonen, H. F. Nielsen and A. Secret, "The World-Wide Web," *Communications of the ACM*, Volume 37, No. 8, August 1994.
- [15] Crowcroft, Jon and Handley, Mark, "Problems with the World-Wide Web," *ConneXions*, Volume 9, No. 6, June 1995.
- [16] Mark Handley and Jon Crowcroft, *The World-Wide Web: Beneath the Surf*, ISBN 1-85728-435-6, UCL Press, 1995.

DAVID STROM is president of his own consulting firm in Port Washington, NY, where he works for leading-edge vendors of networking and communications products. In between consulting gigs, he is publisher, editor, and sole writer of *Web Informant*, a free-for-the-asking HTML-based newsletter distributed via e-mail about Web marketing for the computer industry. His Web site, at URL <http://www.strom.com>, contains back issues along with the usual links to interesting places on the Web that he wishes he thought of first. Strom is a member of the NetWorld+Interop US Program Committee. His background is as an editor and writer in the computer trade press, where he wrote one of the worst articles on the InteropNet ever published. He was founding editor-in-chief of *Network Computing* (the US CMP version) and presently writes for *Infoworld* and *Forbes ASAP* as well as appears occasionally on CBS TV news talking about the Internet, but don't hold any of that against him. E-mail: david@strom.com

What Does the IAB Do, Anyway?

by Brian Carpenter, IAB Chair, CERN

The “Internet Architecture Board” (IAB) sounds as if it is something rather grand, perhaps consisting of a group of people in formal business clothes, sitting around an impressive oak table, under the watchful eyes of an oil painting of The Founder of the Internet. The reality is rather different, and this article is intended to give a feeling for what the IAB really is, what it does, and equally important what it cannot do.

First, the formalities

The IAB was originally called the Internet *Activities* Board, and it was set up in 1983, chaired by Dave Clark, back in the days when the Internet was still largely a research activity of the US Government. The early history of the IAB is hard to trace in detail from the public record, for a reason expressed clearly in the minutes of its meeting in January 1990: “The IAB decided that IAB meeting minutes will be published to the Internet community.” The earlier minutes are not on the public record. A good snapshot of the IAB in 1990, and a short history, are given in RFC 1160, written by Vint Cerf who was the second IAB Chair. He was followed in this post by Lyman Chapin and Christian Huitema. In any case, the 1980s are pre-history as far as the Internet is concerned, and this article concentrates on the present.

Today, the IAB consists of thirteen voting members. Of these, six are nominated each year by a nominating committee drawn from the *Internet Engineering Task Force* (IETF), for a two year term. This membership has to be approved by the Board of Trustees of the *Internet Society*. Indeed, one of the main motivations for the foundation of the Internet Society was to provide a legal umbrella for the IAB and for the IETF’s standardisation actions. The thirteenth voting member of the IAB is the IETF Chair.

In addition, IAB meetings are attended by a representative of the *Internet Assigned Numbers Authority* (IANA) and of the RFC Editor, by a liaison with the *Internet Engineering Steering Group* (IESG), and by the Chair of the *Internet Research Steering Group* (IRSG). Finally, the IAB has a volunteer Executive Director. The IAB elects its own Chair from among its twelve IETF-nominated members.

Now, what are the meetings really like?

Most of the meetings take the form of two-hour telephone conferences about once a month. Due to time-zones, it is early morning for members on the US West Coast, late afternoon for Europeans, and after midnight for our Australian member. Those on the US East Coast have a comfortable mid-morning time slot.

Of course, on a telephone conference, it is hard to see whether the others are wearing smart business suits. In our face-to-face meetings, it’s pretty obvious that most of them are not. These meetings usually take place at IETF locations, 3 times a year. Unfortunately, although we are all in the same time zone physically, it is guaranteed that some of us will be jet-lagged at every meeting. So there is a certain amount of wandering around and a lot of coffee-drinking, but we get through the agenda in the end. An open meeting is also held at each IETF meeting, so that any member of the IETF can address the IAB.

To understand what the IAB really does in its meetings, it is necessary to know that the detailed work of driving the Internet standards process is done by the IESG. Not only must individual members of the IESG, known as *IETF Area Directors*, oversee the work of all the working groups in their area, but the IESG as a group must approve all formal standards actions.

This means approving the conversion of Internet Drafts into Proposed Standards, and subsequent steps towards full standardisation. Since the last set of reforms of IETF process, in 1992–93, the IAB itself does not have to approve individual standards actions.

The IESG consists of a set of specialists in various technical areas, and IESG positions are filled from the IETF by looking for specialists. In contrast, the IAB members are not appointed as specialists, but rather as generalists with good overview of all aspects of the Internet architecture. In a typical meeting, apart from routine business such as reviewing the IAB action list, we will try to discuss one or two strategic issues in some depth. The intention is to reach conclusions that can be passed on as guidance to the IESG, or turned into published statements, or simply passed directly to the relevant IETF working group.

Examples, please!

To give some examples, some issues that have been discussed in recent IAB meetings (those between the July and December 1995 IETF meetings inclusive) were:

- The future of Internet addressing
- Architectural principles of the Internet
- Future goals and directions for the IETF
- Management of Top Level Domains in the Domain Name System
- Registration of MIME types
- International character sets
- Charging for addresses
- Tools needed for renumbering

The IAB does not aim to produce polished technical proposals on such topics. Rather, the intention is to stimulate action by the IESG or within the IETF community that will lead to proposals that meet general consensus. In some cases, the IAB does indeed publish Internet Drafts or RFCs but these are in the nature of statements or viewpoints rather than standards proposals. Past experience has shown that standards proposals that have not passed through the fiery experience of peer review by the IETF are unlikely to be generally accepted.

Another type of action that the IAB can trigger is the setting up of a workshop or ad hoc panel, outside the standards process, to develop ideas in a particular area. For example, workshops were held recently on security (RFC 1636) and information infrastructure (RFC 1862).

The IAB can also stimulate the formation of research groups, which is why the IRSG Chair sits in the IAB. These are expected to have a longer existence than panels or workshops, but do not normally produce standards-track documents.

And in between the meetings?

IAB members try to track the e-mail activity on the main IETF list and on the lists of whichever IETF Working Groups interest them. They can of course intervene as individuals in these discussions whenever they want, but cannot speak in the name of the IAB unless there is a clear consensus.

The IAB Chair, and the nominated IAB Liaison to the IESG, take part in two-weekly IESG telephone conferences and track the e-mail activity of the IESG.

What Does the IAB Do? (*continued*)

While they do not have a vote in formal IESG ballots, they can offer advice on any issue discussed by the IESG and of course refer back to the IAB if necessary. The IAB tends to get involved with IESG discussion at 3 critical junctures in the life of an IETF working group:

- When a new working group is chartered, the IAB may comment on the draft charter before it is approved.
- When documents reach the last call prior to an IESG ballot, IAB members tend to sit up, pay attention, and stick their oar in.
- When a working group gets into a difficult situation, or tension arises between the WG and the relevant IESG Area Director, IAB members try to help individually or collectively to resolve the situation.

Liaison

The IAB also has a role in external representation and formal liaison. The IETF is far from alone in the world of information technology standards. In a few cases (subcommittees of ISO-IEC/JTC1, ITU-T, The ATM Forum), the IETF has established formal liaisons with other bodies, and the IAB (with the Internet Society) has assisted in the bureaucratic part of this. The real technical liaison of course takes place at WG level. More generally, IAB members find themselves contacted by a wide variety of other organisations in search of information, technical contacts, conference speakers, and the like.

According to its charter (RFC 1601), the IAB has several other jobs:

- The IAB appoints the IETF Chair and all other IESG candidates from a list provided by the IETF Nominating Committee. In recent years this has been an easy job, thanks to the excellent job done by the Nominating Committee.
- The IAB provides oversight of the architecture for the protocols and procedures used by the Internet. The IAB has often discussed exactly what this part of its charter means and how to implement it. The activities described above are the practical realisation of this job.
- The IAB provides oversight of the process used to create Internet Standards. The IAB serves as an appeal board for complaints of improper execution of the standards process. Apart from its working relationship with the IESG, IAB members are active in the ongoing review of the current Internet standards process (RFC 1602, under revision). So far there has been only one appeal case, heard in open session by the IAB in April 1995, which resulted in procedural recommendations about the handling of intellectual property rights in the IETF standards process.
- The IAB is responsible for editorial management and publication of the Request for Comments (RFC) document series, and for administration of the various Internet assigned numbers. In fact these responsibilities are fully delegated to the RFC Editor and the IANA respectively, with the IAB available for consultation when needed.
- The IAB acts as a source of advice and guidance to the Board of Trustees and Officers of the Internet Society concerning technical, architectural, procedural and (where appropriate) policy matters pertaining to the Internet and its enabling technologies. It must be said that this channel has been little used, and a more regular contact between the IAB and the Internet Society is highly desirable.

What we do not do

The IETF is a standards body and the IAB is drawn from the IETF in order to help it achieve its goals of better standardisation. For this reason, the IAB has no official role in operational or commercial matters and only a minor role in policy matters. As an example, the IAB could decide to stimulate work on a standard for automatic labeling of e-mail describing how to build nuclear warheads, but could not make policy on whether such messages should be forbidden on international routes.

However, the boundaries of the proper role for the IETF, the IESG and the IAB are somewhat fuzzy. A real-life example was the request from the Internet Society to the IAB for work to be done on Internet ethics. Although this was hardly standards work, the IETF responded by setting up the working group called "Responsible Use of the Network," which recently finalised RFC 1855, the Netiquette guide.

At the time of writing, the IAB is involved in a discussion about the future management of the Top Level Domains of the DNS. Given that this is an existing responsibility of the IANA, any proposal for change automatically involves the IAB, and raises another example of the ambiguous boundary between standards, policy and operations.

Another fuzzy boundary is "how far up or down do we go?" With the international political drive for information superhighways, the IAB is expecting the Internet to become the infrastructure for the "Information Infrastructure." Does this mean that every information handling protocol must be developed by the IETF? Certainly not! So there will be a boundary between the IETF standards and other information handling standards, and this will not be a completely clear boundary. Similarly, the boundary between IETF standardisation and hardware transmission standardisation can never be rigid. This is particularly apparent in the case of ATM. The IAB has a role to play in defining the limits of the IETF, closely linked to the question of liaison with other bodies.

In conclusion

The IAB exists to serve and help the IETF, attempting to strike a balance between action and reaction. IAB members are part-time volunteers (as indeed are IESG members) serving the IETF community with no particular expectation of reward. Of the 12 nominated IAB members at the time of writing, seven are based in the USA, with one each in Australia, Canada, France, the Netherlands and Switzerland. Eight of us work in the computing and telecommunications industry, one in manufacturing industry, two for government-funded research institutes, and one for a university. The days when the IAB could be regarded as closed body dominated by representatives of the United States Government are long gone. Any IETF member can volunteer for the Nominating Committee, in order to influence the future membership of the IAB.

More information

The IAB has a Web server with URL <http://www.isi.edu/iab/>. IAB minutes are also accessible using anonymous FTP from host [ftp.isi.edu](ftp://ftp.isi.edu/pub/IAB), in directory `pub/IAB`. General information about the IETF is available with the URL <http://www.ietf.org/home.html>, or from RFC 1718 ("The Tao of the IETF").

BRIAN E. CARPENTER has been Group Leader of the Communications Systems group at CERN since 1985, following ten years' experience in software for process control systems at CERN, which was interrupted by three years teaching undergraduate computer science at Massey University in New Zealand. He holds a first degree in physics and a Ph.D. in computer science, and is an M.I.E.E. He is Chair of the Internet Architecture Board and an active participant in the Internet Engineering Task Force. E-mail: brian@dxcoms.cern.ch

To Foster Residential Area BroadBand Internet Technology

IP Datagrams Keep Going, and Going, and Going...

by Mark Laubach Com21, Inc.

Abstract

Cable modem technology is rapidly entering commonplace discussion. The capabilities provided by cable modems promise data bandwidth speeds far in excess of those provided by traditional twisted pair public telephone networks. Internet service providers are taking position to promote this next generation method of delivering Internet services to the home as part of the broadband access to the home race. Cable TV operators and Regional Bell Operating Companies (RBOCs), e.g., Pac*Bell, are preparing for this integrated broadband future by installing or rebuilding existing all-coaxial cable plants into two-way *Hybrid-Fiber Coaxial* (HFC) plants, and by offering a wide range of both data and interactive services which they feel will be most attractive to their subscriber base. Initially these services will only provide Internet access and access to the major information service (e.g., CompuServe, AOL, and Prodigy). These service offerings will quickly advance to support multi-player gaming and collaborative services such as voice and desktop video teleconferencing.

As an introduction to some of the issues surrounding cable modem technology, this article summarizes two of the standardization efforts: the ATM over HFC definition work taking place in the ATM Forum's Residential Broadband Working Group, and the standards progress in the IEEE P802.14 Cable TV Media Access Control and Physical Protocol Working Group. Delivering a viable Internet service to a Cable TV based subscriber community has its own set of deployment issues that are briefly overviewed and summarized.

Introduction

Packet technology has been around since 1964 [8]. Since then, the size of packets has been debated, as well as variable versus fixed size. Packets are transmitted over any media these days however, the next economic and technical frontier is moving packets over Cable TV (CATV) networks. The driving push is to deliver IP datagrams over CATV networks. There are several datalink methodologies for delivering IP datagrams via cable modems. This article overviews the notion of sending small, fixed sized packets over the CATV plant. These packets are 53-octet *Asynchronous Transfer Mode* (ATM) cells [1].

Numerous standards organizations are gearing towards producing cable modem standards. The ultimate goal of each is to drive cable modem availability to commodity status and made available via consumer "off the shelf" purchases at computer boutiques and electronic supermarkets. The minor problem with the commodity process is that these numerous standardization activities are competing and largely uncoordinated and there are about a dozen cable modem manufacturers producing product, some of whom wish to establish de facto standard status by being first to market.

The IEEE P802.14 Cable TV MAC and PHY Protocol Working Group is chartered with providing a single MAC and multiple PHY standard for cable TV networks. P802.14 must support IEEE 802 layer services (including Ethernet) and must also be "ATM Friendly." ATM residential broadband work is currently taking place in the ATM Forum.

The customer network interface "du jour" is Ethernet 10Base-T. There is a mandate for a 10Mbps Ethernet interface in the home. Subscriber access equipment can be a personal computer, X-Terminal, or any such device which support the TCP/IP protocol suite.

Engineering challenges of data over cable

The standardization and implementation of two-way interactive services on Hybrid Fiber-Coax (HFC) TV networks is fraught with many engineering problems which must be overcome: 1) cable TV systems are inherently asymmetric in nature, i.e., there is more downstream bandwidth available than upstream and interactive services such as voice or video telephony require symmetric data rates; 2) the upstream facility is typically located in a sub- or low-split frequency region, typically from 5 Mhz to 40 Mhz, which is laden with many noise impairment sources and is termed a "hostile RF environment" by many organizations who have taken a close look at the situation [2, 11]; 3) high utilization of the upstream bandwidth is necessary and accomplished by sharing bandwidth between stations with the access based on dynamic assignments within a slotted regimentation; 4) current IEEE standard efforts are looking towards supporting a 50 mile cable length which places approximately 400 microseconds of propagation delay between the cable head-end and the farthest subscriber station out on the network—this requires precise ranging of subscriber stations to support a high utilization sharing architecture on the upstream channel, and 5) the presence of impulse noise, ingress noise, common-mode noise, micro-reflections, group delay, etc., and the effect of inexpensive "do it yourself" CATV home wiring create the upstream impairment challenge and require a robust modem coding mechanism. The industry challenge is to provide robust upstream RF modem that can cope with these impairment problems. Solutions include the use of coding techniques that can operate in a lower signal to noise environment and by the use of techniques such as *Forward Error Correction* (FEC).

Once the robust RF Physical (PHY) modem environment is in place, a Media Access Control (MAC) protocol layer can be layered or coupled to the PHY creating a means to pass data-link layer information between cooperating stations. The choice of the allocation protocol and the placement of the bandwidth ownership intelligence is important. A straightforward approach is to place the ownership of the upstream bandwidth under the direction of the head-end controller. This also has the effect of reducing complexity in the subscriber unit and by centralizing the allocation intelligence in the network. Communications between the head-end controller and each subscriber unit is important as permission to use the upstream channel is granted by the head-end controller whose allocation algorithm must take into account the needs communicated to it by each subscriber unit.

The system architecture must be constructed such that both small and large systems can be built that work in the variety of cable television plants that exist today, and subscriber system must be easily installed and the whole system made manageable and scalable.

The selection of ATM cells as the data-link layer protocol data unit for Cable TV networks has the advantage in that it provides a suitable integrated multiplexing platform capable of supporting a mix of guaranteed (predictive) traffic flows with best-effort (reactive) traffic flows. In addition, the nature of ATM allows other multimedia applications to be added in the future without requiring iterative changes to the basic ATM protocol. Cable operators can deploy ATM systems as part of an evolutionary path to a fully integrated multimedia bearer service offering.

An ATM data-link protocol can be layered in a straightforward manner for both the downstream and upstream segments of a cable modem network.

ATM in the Residential Broadband Network

Residential Broadband Technology (continued)

The challenges are that upstream traffic management and resource management must be creatively controlled to support the guaranteed and best effort *Quality of Service* (QoS) needs of the cable modem. A residential ATM bearer service easily supports Internet access to the home via the Classical IP over ATM standards of the Internet Engineering Task Force (IETF) [3] or by providing an IP over Ethernet adaptation overlay service.

While ATM in the home is desired as a future interconnection method by some HFC operators, the cost burden of the ATM interface is not economically feasible today. It is expected that ATM network interface controllers will be decreasing in cost quickly over time so planning a cable modem bearer service now to support both Ethernet and ATM home interfaces can be viewed by some as prudent.

GENERAL:

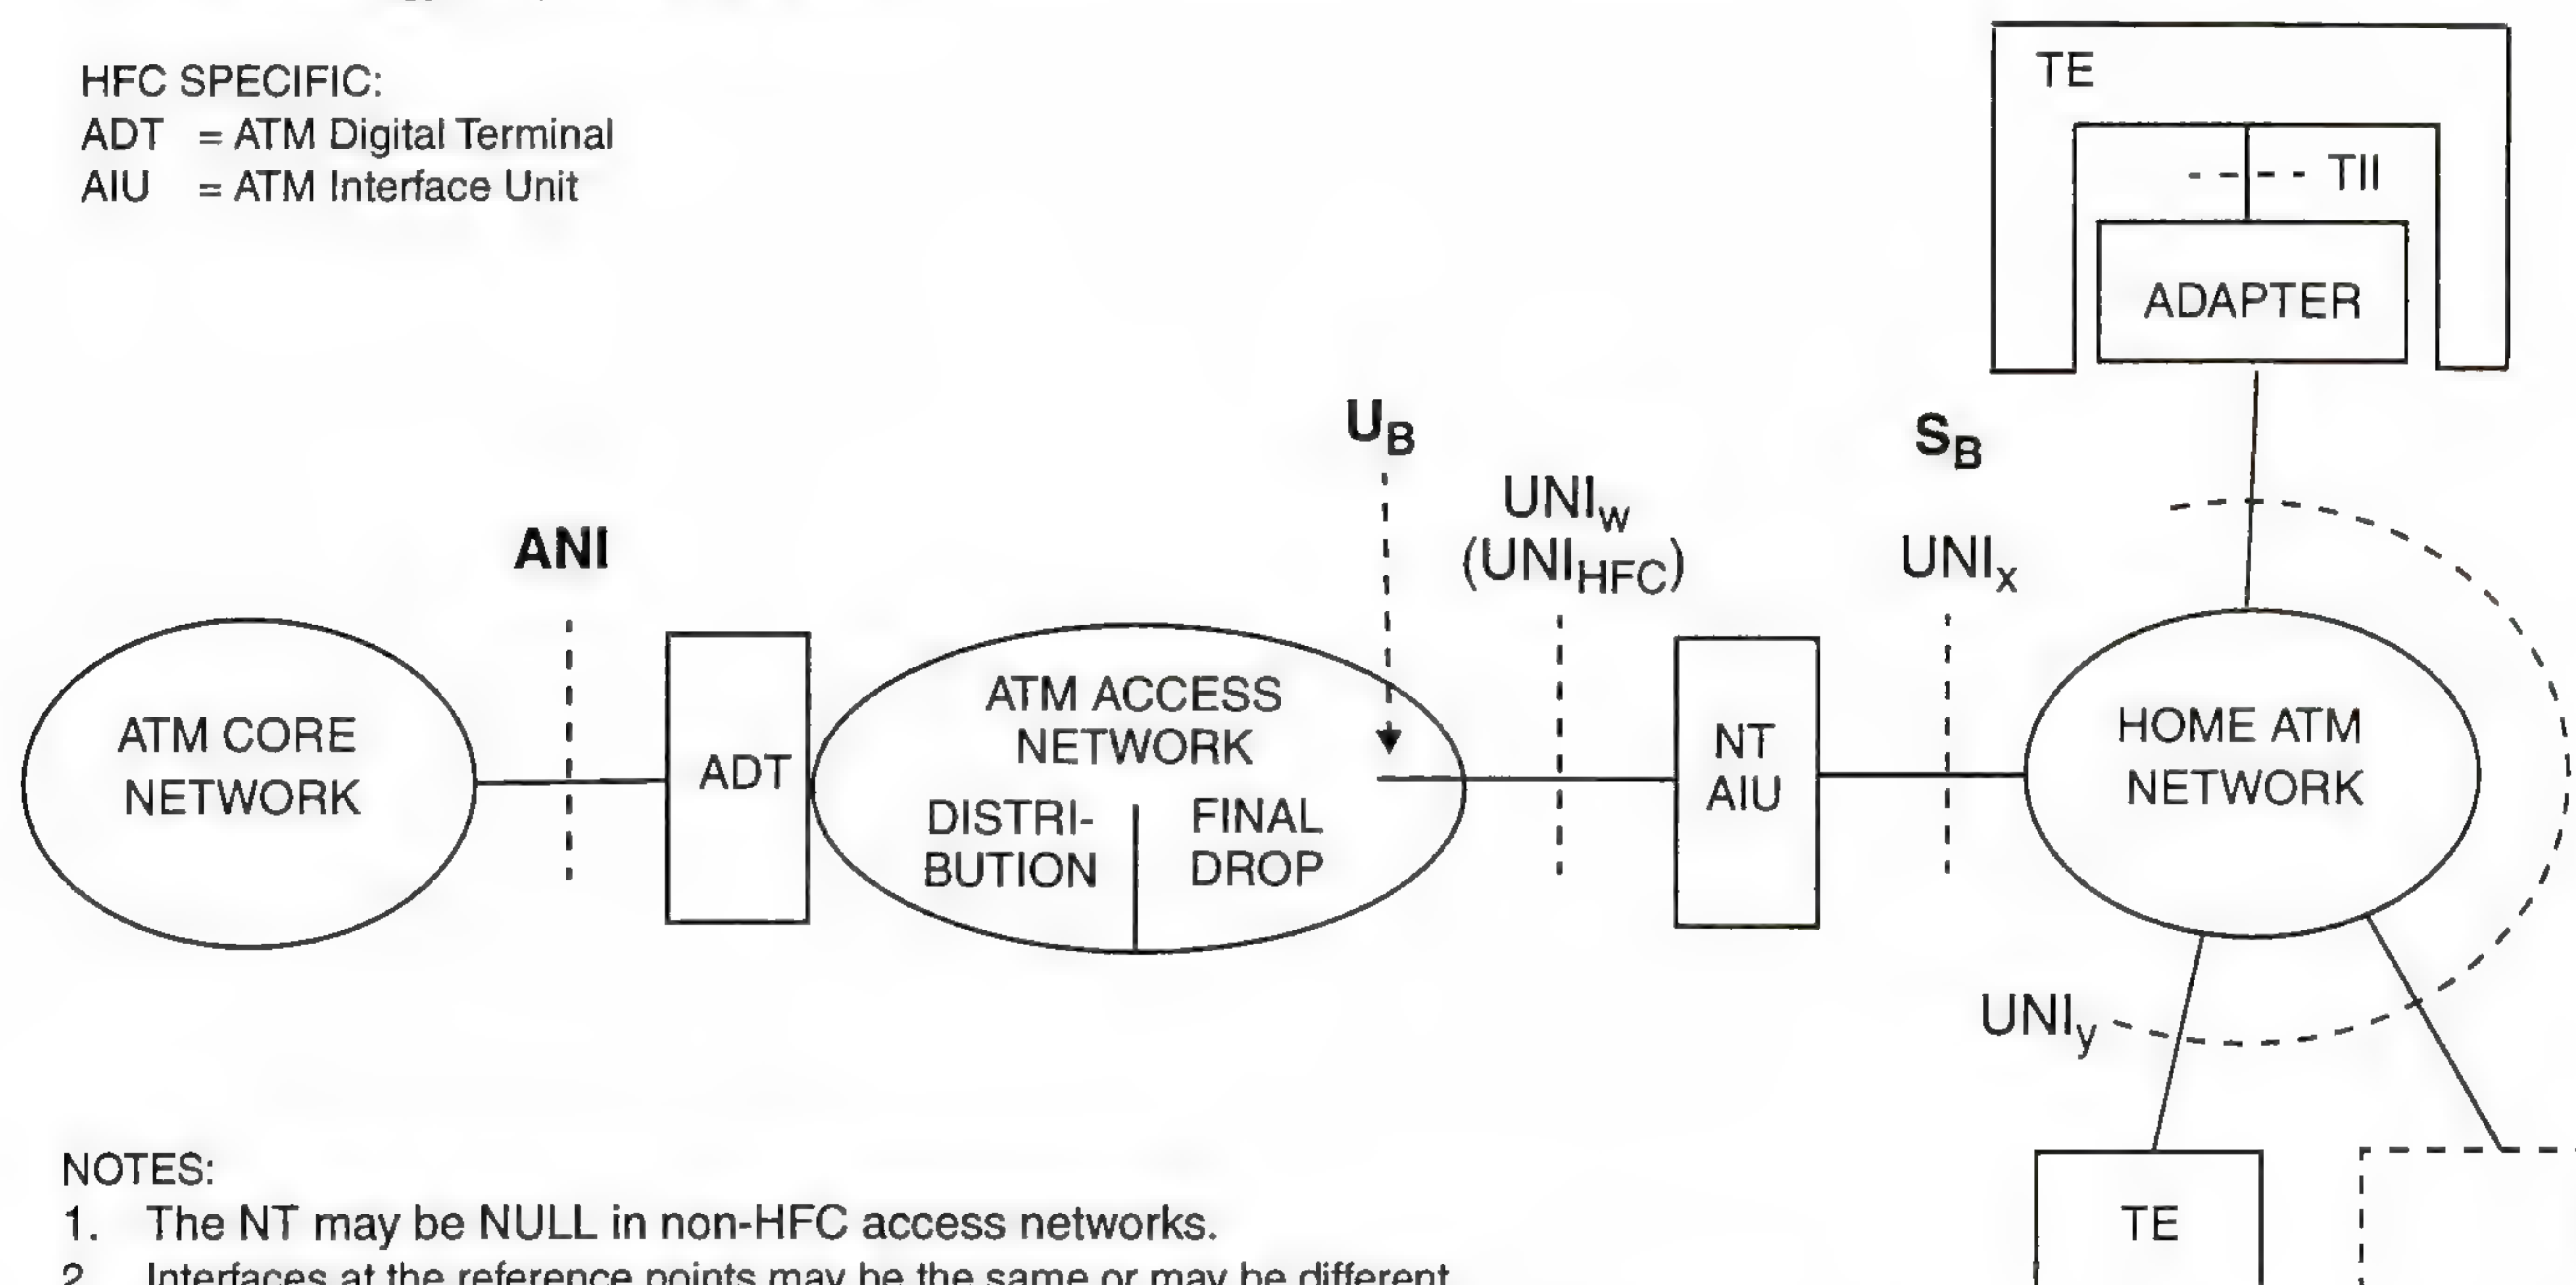
ANI = Access Network Interface

TII = Technology Independent Interface

HFC SPECIFIC:

ADT = ATM Digital Terminal

AIU = ATM Interface Unit



NOTES:

1. The NT may be NULL in non-HFC access networks.
2. Interfaces at the reference points may be the same or may be different.
3. The HAN contains physical media and passive devices. It may also contain active devices; e.g. bridges or switches
4. There will be more than one PMD and MAC layer specified for the UNI reference points

Work in Progress - Based on RBB Reference Configuration - August 1995 - Version 3.0

Figure 1: ATM Forum Residential Broadband Reference Model

The ATM Forum's RBB Working Group

The ATM Forum is focusing attention on delivering ATM over residential broadband distribution systems. This work is being carried out in the *Residential Broadband* (RBB) Working Group (WG). The material presented in this section represents work in progress in the ATM Forum and is offered as an example of the current thinking on the subject. At some time in the future, the ATM Forum will be producing a published specification which includes the ATM over HFC UNI details. The ATM Forum is a closed industrial consortium requiring membership dues for participation.

The two goals of the RBB WG are to 1) deliver ATM to the home and 2) deliver ATM within the home. These can be euphemistically termed as “the last mile” problem and “the last yard” problem. The current proponents of ATM over HFC systems are concerned with delivering a full function UNI interface to the home via an active *Network Interface Unit* (NIU) termed an *ATM Interface Unit* (AIU). Controlling the system is an *ATM Digital Terminal* (ADT) located at the cable system head-end (see Figure 1). The discussion of ATM within the home is beyond the scope of this article.

The *ATM Network Interface* (ANI) defines the connection between the ADT and the ATM WAN network. This interface may either be specified as a *Network–Network Interface* (NNI) or as a UNI. The ANI will be based on existing ATM standards and the WG expects complete compliance with existing physical (PHY) interface standards.

The HFC access network is in effect a “black box” to the ATM Forum’s design activities. It was decided early in the RBB charter process, that the RBB WG will rely on the efforts from IEEE P802.14 Cable TV MAC and PHY Working Group for the transport of ATM cells over the HFC network. The UNI_{HFC} will define an RF interface for the ADT and AIU. A possible protocol stack representation of the relationship between ATM and P802.14 is shown in Figure 2. It has been suggested that communication between the ATM layer and the IEEE access layer be specified via an abstract layer interface definition, which will be referred to as the *access interface* in this article.

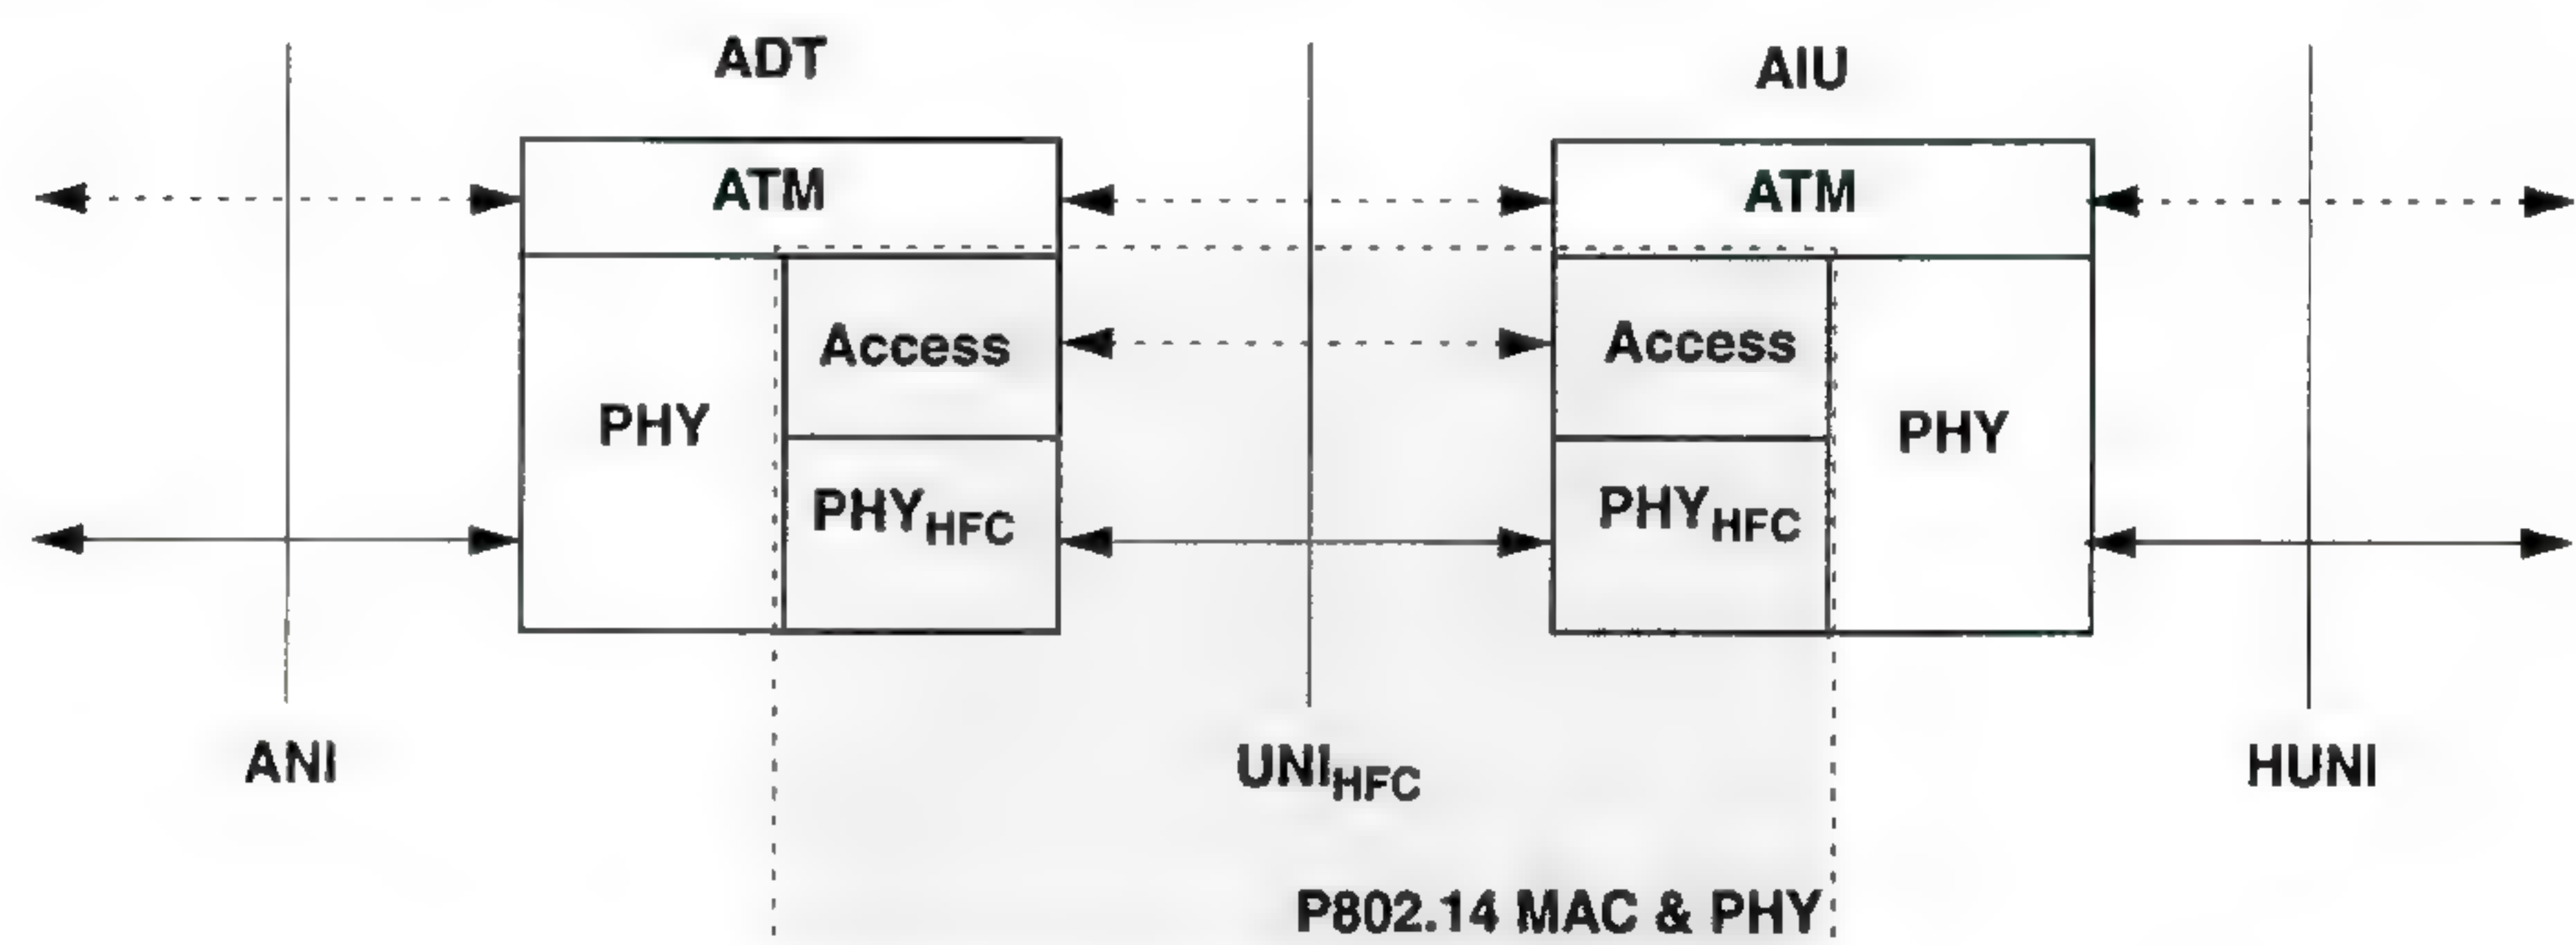


Figure 2: Proposed ATM Transport Protocol Model

The AIU provides an ATM UNI to the home. This *Home UNI* (HUNI) interface is meant to be as standard as possible. It will most likely provide a subset of UNI 3.1 or the upcoming UNI 4.0. It is expected that this interface will deliver the full range of CBR, VBR, ABR, and UBR services. However, the real performance of these will be limited by the available performance offered by the HFC access network and the characteristics of the underlying P802.14 MAC and PHY.

Issues

The RBB WG effort is current work in progress. It will produce an implementation reference (i.e., a UNI_{HFC} implementation reference) which is synchronized with the IEEE P802.14 working group. At the time of this writing, the following issues will need to be resolved within the RBB:

- How Virtual Path Identifiers (VPIs) and Virtual Circuit Identifiers (VCIs) will be used within the UNI_{HFC} and how they are mapped to the IEEE access interface.
- Where and how does ATM UNI Traffic Management (TM) take place in the HFC system, and the nature of the QoS or TM interface to the IEEE access interface.

continued on next page

Residential Broadband Technology (*continued*)

- What form of UNI signaling will be supported between the ADT and the AIU? The AIU might be passive requiring the ADT to perform proxy signaling on behalf of the home UNI. If all AIU interfaces share a common VCI space, then meta-signaling may be required, etc.
- Will the ATM over HFC system specification include telephony voice over ATM and if so, with what Cell Delay Variation (CDV)?
- What are the required performance goals for ATM peer-to-peer networking when operating over an HFC network.?

The above issues and more will be addressed in the ATM Forum's efforts.

IEEE P802.14 Cable TV MAC and PHY Protocol Working Group

In November 1994 the IEEE P802.14 CATV MAC and PHY Protocol Working Group met for the first time as an approved project within the 802 standards committee. Previous work had been done in the 802.catv study group in preparation for formal approval. The Project Authorization Request (PAR) charter of the group specifies that it will standardize a single MAC layer protocol and multiple PHY layer protocols for two-way HFC networks. Consistent with the IEEE LAN/MAN 802 Reference Model [5], P802.14 will produce a solution which supports the 802 protocol stack while at the same time supporting ATM in an "ATM Friendly" manner.

The WG has completed a first release revision of a functional requirements document [4] which details the P802.14 cable topology model (see Figure 3); defines key assumptions, constraints, and parameters; defines key performance metrics and criteria for the selection of multiple PHY protocols and a MAC protocol; and defines the support of QoS parameters. The WG's work plan called for the close of formal proposals in November 1995, with the recommended protocol defined in July 1996. Seventeen MAC protocol proposals were submitted to the working group. It is anticipated that the WG will select the best features from amongst the proposals and apply appropriate glue to form the standard. The IEEE is a public standards organization and anyone may participate in the standards activities.

The branch and tree topology of a Cable TV single-cable distribution network (see Figure 3 and 4) is divided by RF frequency into a downstream portion (typically 50Mhz to 550Mhz or 750Mhz) and an upstream portion (typically 5Mhz to 40Mhz). Both downstream and upstream frequencies are active on the same physical coaxial RF cable, and the use of bandpass filters and diplexors provide the spectral separation necessary for the simultaneous amplification of signals in each direction. A P802.14 subnetwork can be thought of as a single head-end controller communicating with a set of cable modems via a MAC protocol operating on a collection of downstream and upstream PHY channels (see Figure 4).

Home Network Interface Units (NIUs) will communicate with the head-end terminal unit using the agreed upon downstream and upstream PHY. The downstream PHY will support a broadcast, one transmitter-many receiver model. The upstream PHY will support a one receiver-many transmitter model, requiring that the upstream PHY be shared amongst all participating NIUs in the subnetwork.

Requirements

The general P802.14 requirements are: support of symmetrical and asymmetrical rates on connections involving the downstream and upstream channels; support of *Operation, Administrations, and Maintenance* (OAM) functions; support of one way delays on the order of 400 microseconds (round trip delays to 800 microseconds) support of a large number of users; support for moving data from an originating subnetwork to a destination subnetwork which may be the same or a different one; and the option of a customer reference point between in-home and external networks.

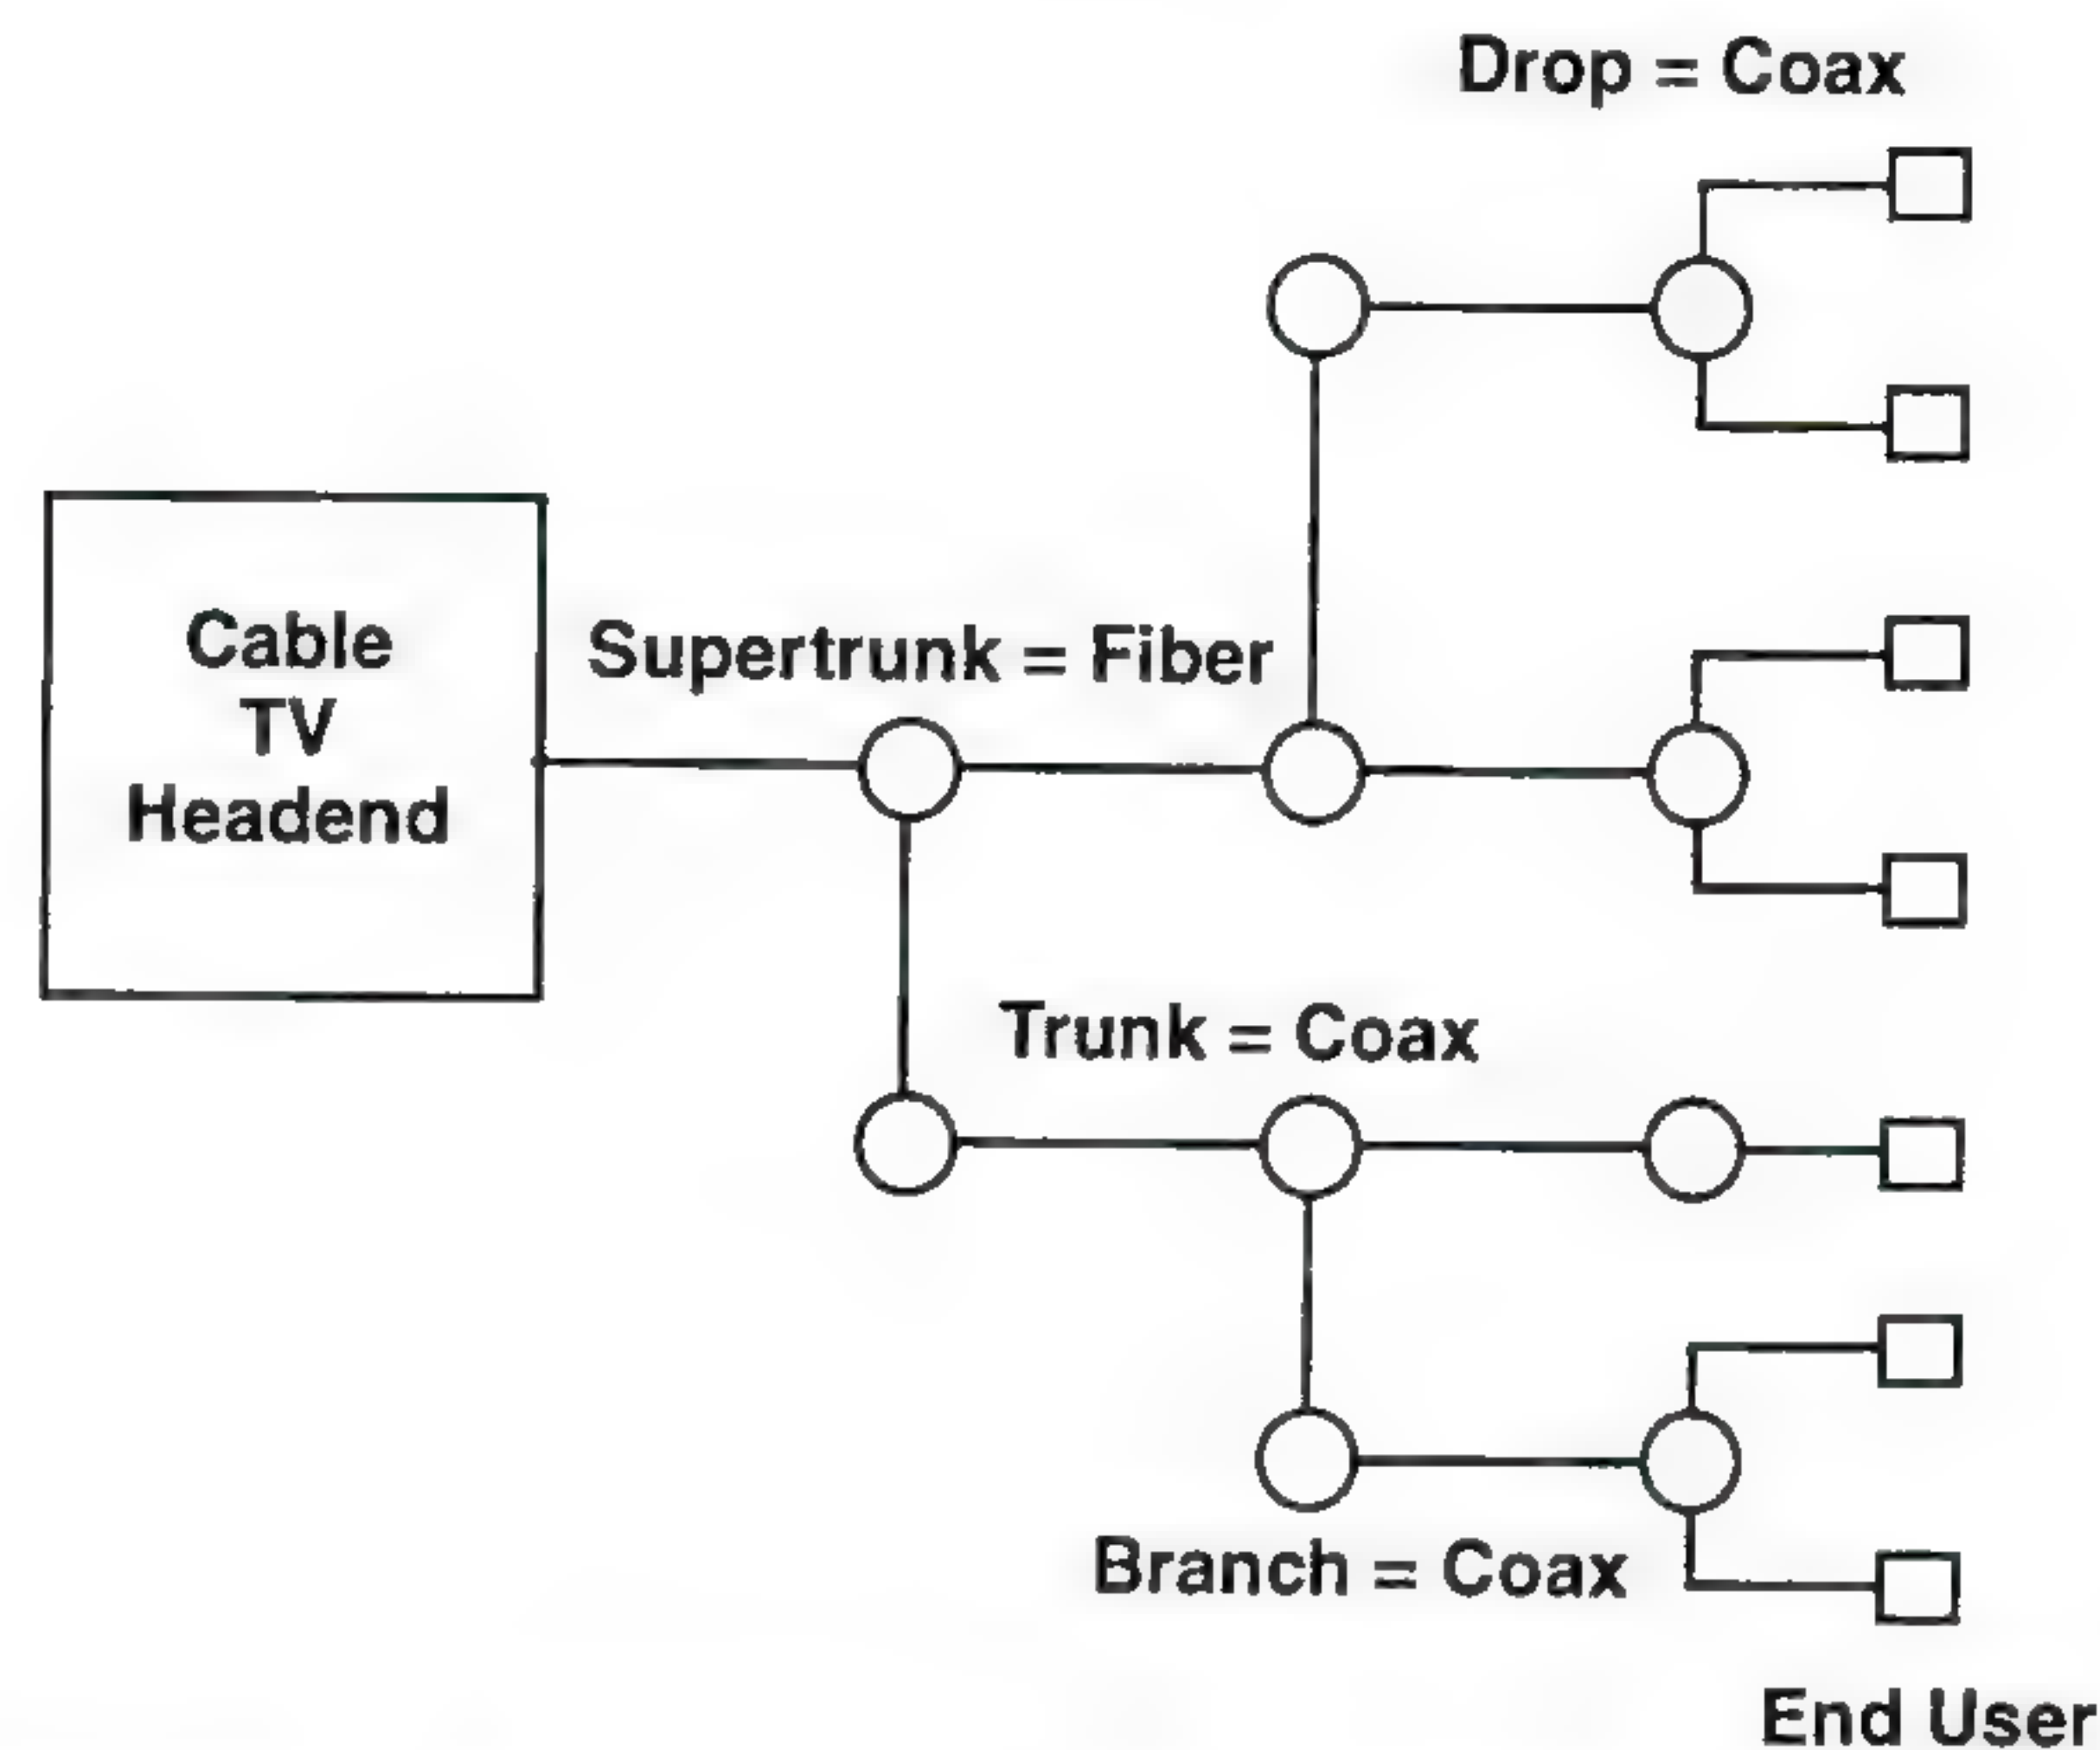


Figure 3: IEEE P802.14 HFC System Topology

The P802.14 MAC layer requirements are: support of both connectionless and connection-oriented services; support of a formal QoS for connections; support for dynamically allocated bandwidth for different types of traffic, including *Constant Bit Rate* (CBR), *Variable Bit Rate* (VBR), and *Available Bit Rate* (ABR); support for unicast, multicast, and broadcast services; interoperability with ATM; predictable low average access delay without sacrificing network throughput; and fair arbitration for shared access to the network within any level of service.

The P802.14 PHY layer requirements are: HFC system size up to 500 households as a reference design point; primary support of sub-split cable plants (5Mhz to 40Mhz upstream), with optional support of mid-split (5Mhz to ~120Mhz upstream) and high-split (~800Mhz to ~1GHz upstream); frequency reuse in the upstream channel; and co-existence with other home information appliances (e.g., entertainment TV) and other uses sharing the broadband system.

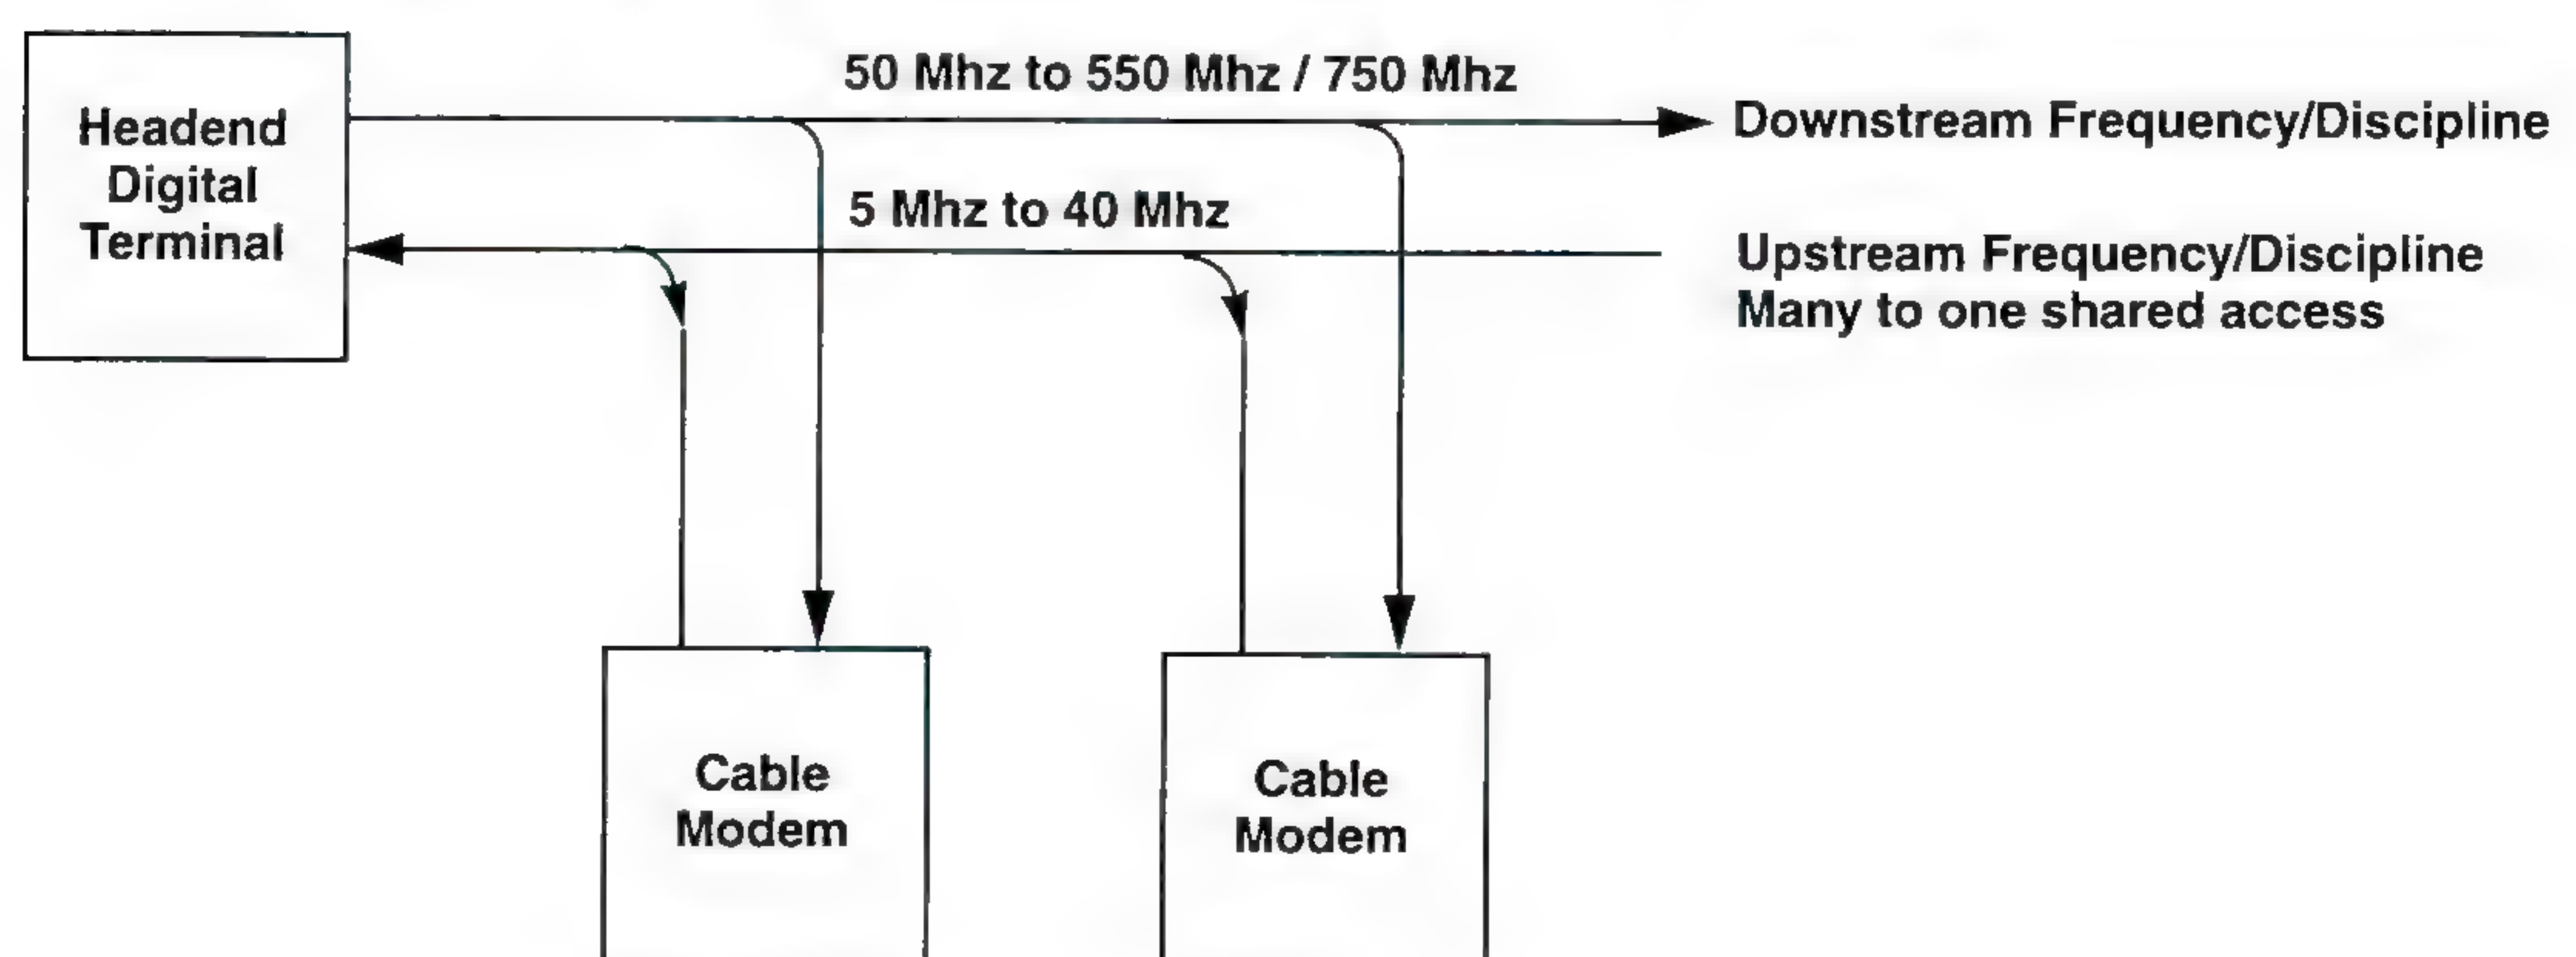


Figure 4: IEEE P802.14 Shared HFC Architecture

continued on next page

Residential Broadband Technology (continued)

The detailed performance requirements for the MAC and for the PHY have yet to be specified by the P802.14 WG. The majority of the MAC proposals received by the working group will be put to modeling and simulation performance scrutiny with the initial results presented at the March 1996 meeting.

There has been much discussion within the P802.14 WG as to the protocol stack model to be used in the standard. One possibility is to treat the 802.2 Logical Link Layer as a peer with the ATM layer, with each interfacing to the 802.14 MAC layer via an access interface, which in turn uses the 802.14 PHY. This is the ATM Friendly model. The other model that has been suggested is layer all 802 MAC layers over an ATM and ATM Adaptation Layer (AAL) stack [1], which in turn uses the 802.14 PHY. This is the All ATM model. These two approaches are shown in Figure 5. While it is premature to say which answer will be the official one, it is useful to note that the ATM Friendly model and its access interface is consistent with both the ATM Forum’s model and the IEEE 802 LAN model.

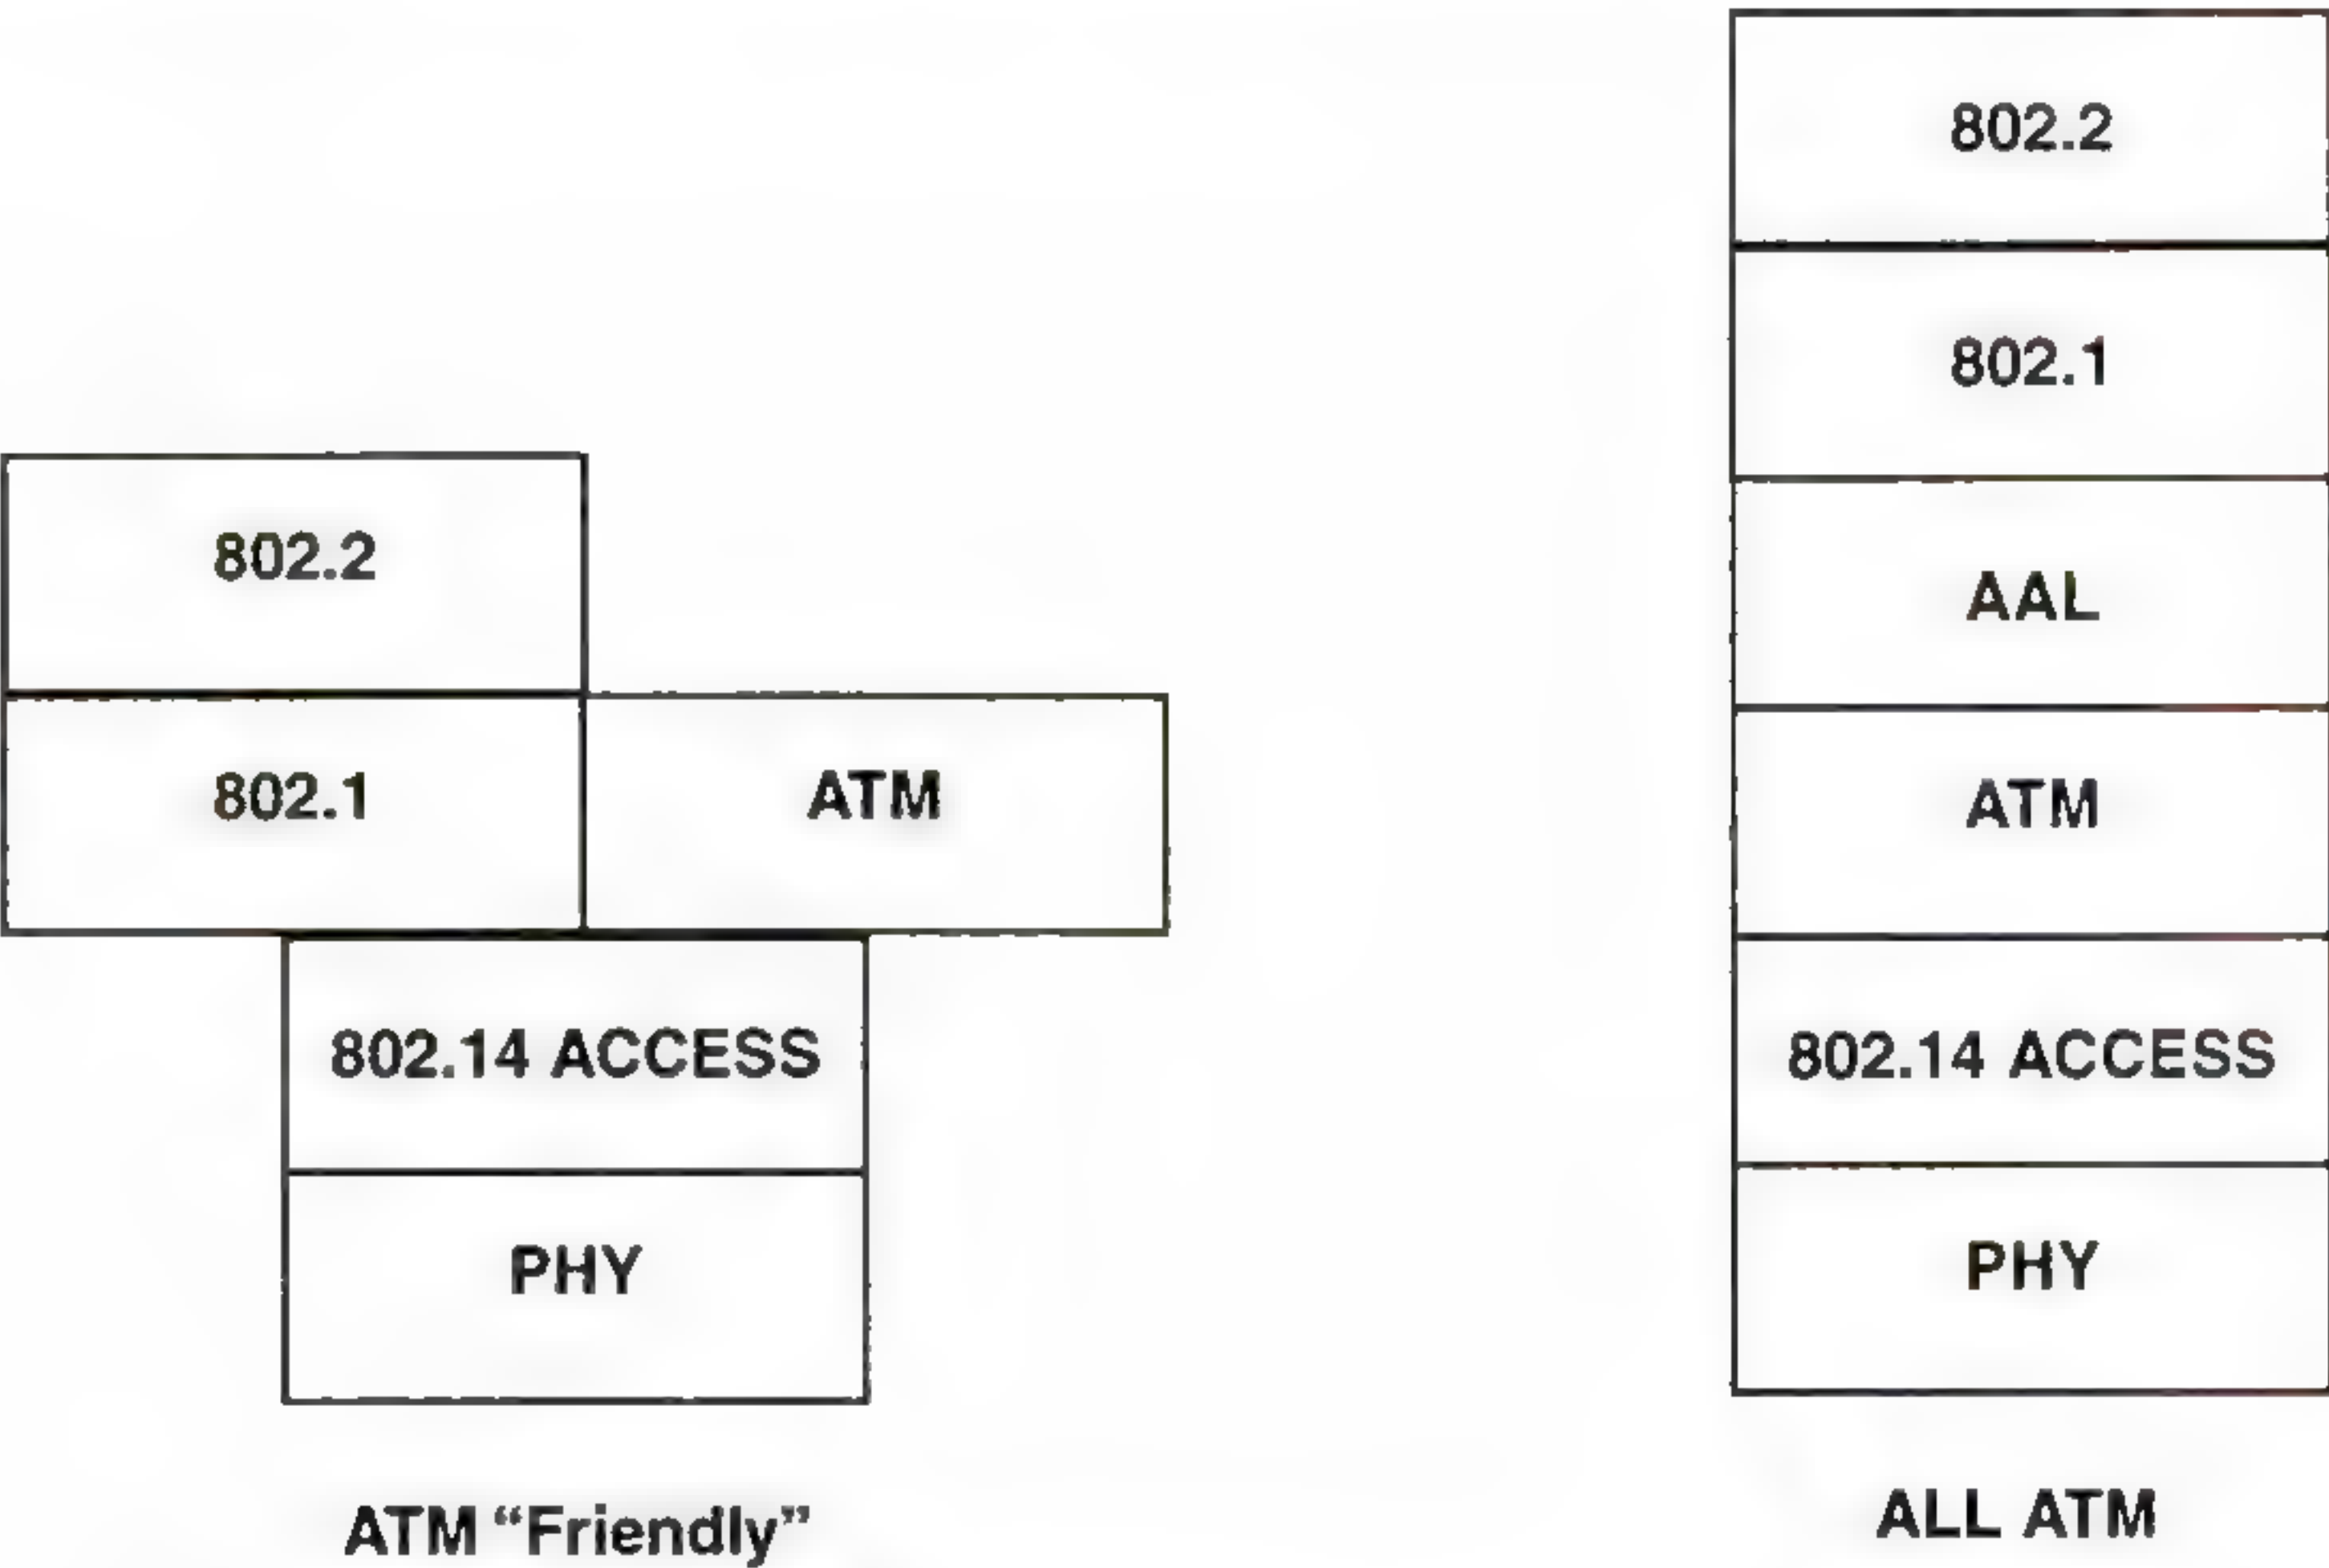


Figure 5: IEEE P802.14 Protocol Stack Alternatives

- Issues** The P802.14 work plan has been recently finalized, setting the stage for a completed draft work by the end of 1996. This section has attempted to summarize some of the aspects of the challenge faced by the P802.14 WG. At the time of this writing, the following additional issues will need to be resolved by the WG:
- Will Plain Old Telephone (POTS) be a fundamental service; i.e., DS0 at 64 Kbps?
 - Will ATM be selected as the Protocol Data Unit (PDU)?
 - What Forward Error Correction (FEC) algorithm will be used and how much protection?
 - If a slotted approach is used, what is the size of the slots?
 - Where will complexity be placed in the system? Putting it where is easiest to fix/maintain implies the head-end.
 - Many-to-one sharing of a single upstream channel using a slotted approach requires ranging of the home NIUs. How precise will the ranging be and how will it be performed?
 - Will the WG specify a set of PHY profiles that may be used or just one downstream and one upstream PHY?

- How will provisioning of authorized stations be performed by the MAC protocol?
- How will the downstream and/or upstream channels be encrypted and how will keys be managed?
- How does the MAC protocol handle errors?
- How will stations will identified in the subnetwork?

The above issues are continually being discussed in the P802.14 working group. As of this writing, the P802.14 Working Group has tentatively decided to select *Quadrature Amplitude Modulation* (QAM) 64 as a mandatory protocol for the downstream channel. The use of QAM16 and QAM256 are for future study. The modulation technique for the upstream channel is currently under debate in the working group. There is a reasonably likelihood that *Quadrature Phase Shift Keying* (QPSK) will be selected due to its ability to perform better in a low signal to noise environment and that there is past industrial experience with this method. The specific choice will be made during the March 1996 meeting.

Expected downstream and upstream data channel rates

For downstream, the QAM64 technique is a 6 bit per Hertz coding scheme which yields a raw data rate of approximately 30Mbps in a North American 6 Mhz wide standard video channel (data + guard bands). With FEC and the effects of framing, the actual usable information data rate is approximately 27Mbps. The usable bandwidth is shared amongst all cable modems for both user and management traffic.

Due to the noisy conditions of the upstream cable environment, the modulation technique which will mostly likely be selected is QPSK. The raw channel rate will be anywhere from 1.5 to 3.0 Mbps with the specific rate selected in the near future by the IEEE P802.14 Working Group. This article will use an example rate of 2.56Mbps raw within a 1.8Mhz bandwidth allocation. The upstream channel requires a longer preamble and more FEC as compared to the downstream channel. The requirements for sharing an upstream channel between multiple modems mandates the use of a guardband (dead time) between packet bursts. The information data rate of a single upstream channel will be approximately 2.0Mbps. It is expected that several upstream channels can be used with the single downstream channel. The head-end controller will "place" cable modems on the appropriate upstream channels to facilitate load balancing and robustness needs.

An IP over cable modem example

This section presents a brief overview of a hypothetical IP over HFC system. It is meant as an informative example designed to illustrate the application of the IP technology and some of the issues surrounding providing the service over a residential cable TV network. Moving IP datagrams in and out of the home over the cable plant is the important issue. The specific technology and protocols used by the cable modem vendor is important only for its ability to provide the required IP service support.

For this example, consider a system that has the following design goals and requirements:

- One-to-many will be supported on the downstream; that is many cable modems are reachable via the downstream channel.
- Many-to-one will be supported on the upstream, i.e., sharing of the upstream channel bandwidth. There may be up to several upstream channels.

continued on next page

Residential Broadband Technology *(continued)*

- The protocol used between the head-end controller and the head-ends is not significant so long as it meets the needs of the IP service.
- The head-end owns the upstream bandwidth and allocates resources to cable modems
- IP over Ethernet 10Base-T is the required interface in the home.
- IP over Ethernet or IP over ATM is the required interface at the head-end.

This example will rely on the ATM Forum and IEEE P802.14 information presented previously in this article. In the downstream channel, a P802.14 destination address will be required for the ATM Friendly model as a unicast, point-to-point VCI will need to be sent to a particular cable modem amongst all the cable modems listening to the downstream. Broadcast traffic can use a special value of the destination address to indicate an “all cable modems receive” message. Special values of station identifiers can be used for multicast groups.

The head-end controller can transmit cells to any cable modem on the channel in any order or rate appropriate to the scheduling information it has and controls. The controller should also participate in the *IP Multicast Group Membership Protocol* (IGMP) and the *IP Resource ReserVation Protocol* (RVSP) and make changes in the cable modem resource assignments and allocations as needed. The home cable modem is only permitted to use the upstream channel under direction of the head-end controller. Guaranteed and best effort bandwidth allocations are dynamically assignable by the head-end controller. It is assumed that the cable modem has a bandwidth request facility for either asking the head-end controller for bandwidth or for communicating its queue status (e.g., “please give me ten seconds at 64Kbps” versus “there are 10 packets in my queue”). The function of the bandwidth management process is to sort these requests for service and give fair access to the requesting cable modems. The specific requirements of the bandwidth manager will be worked out and specified in detail in the ATM Forum RBB and the IEEE P802.14 efforts.

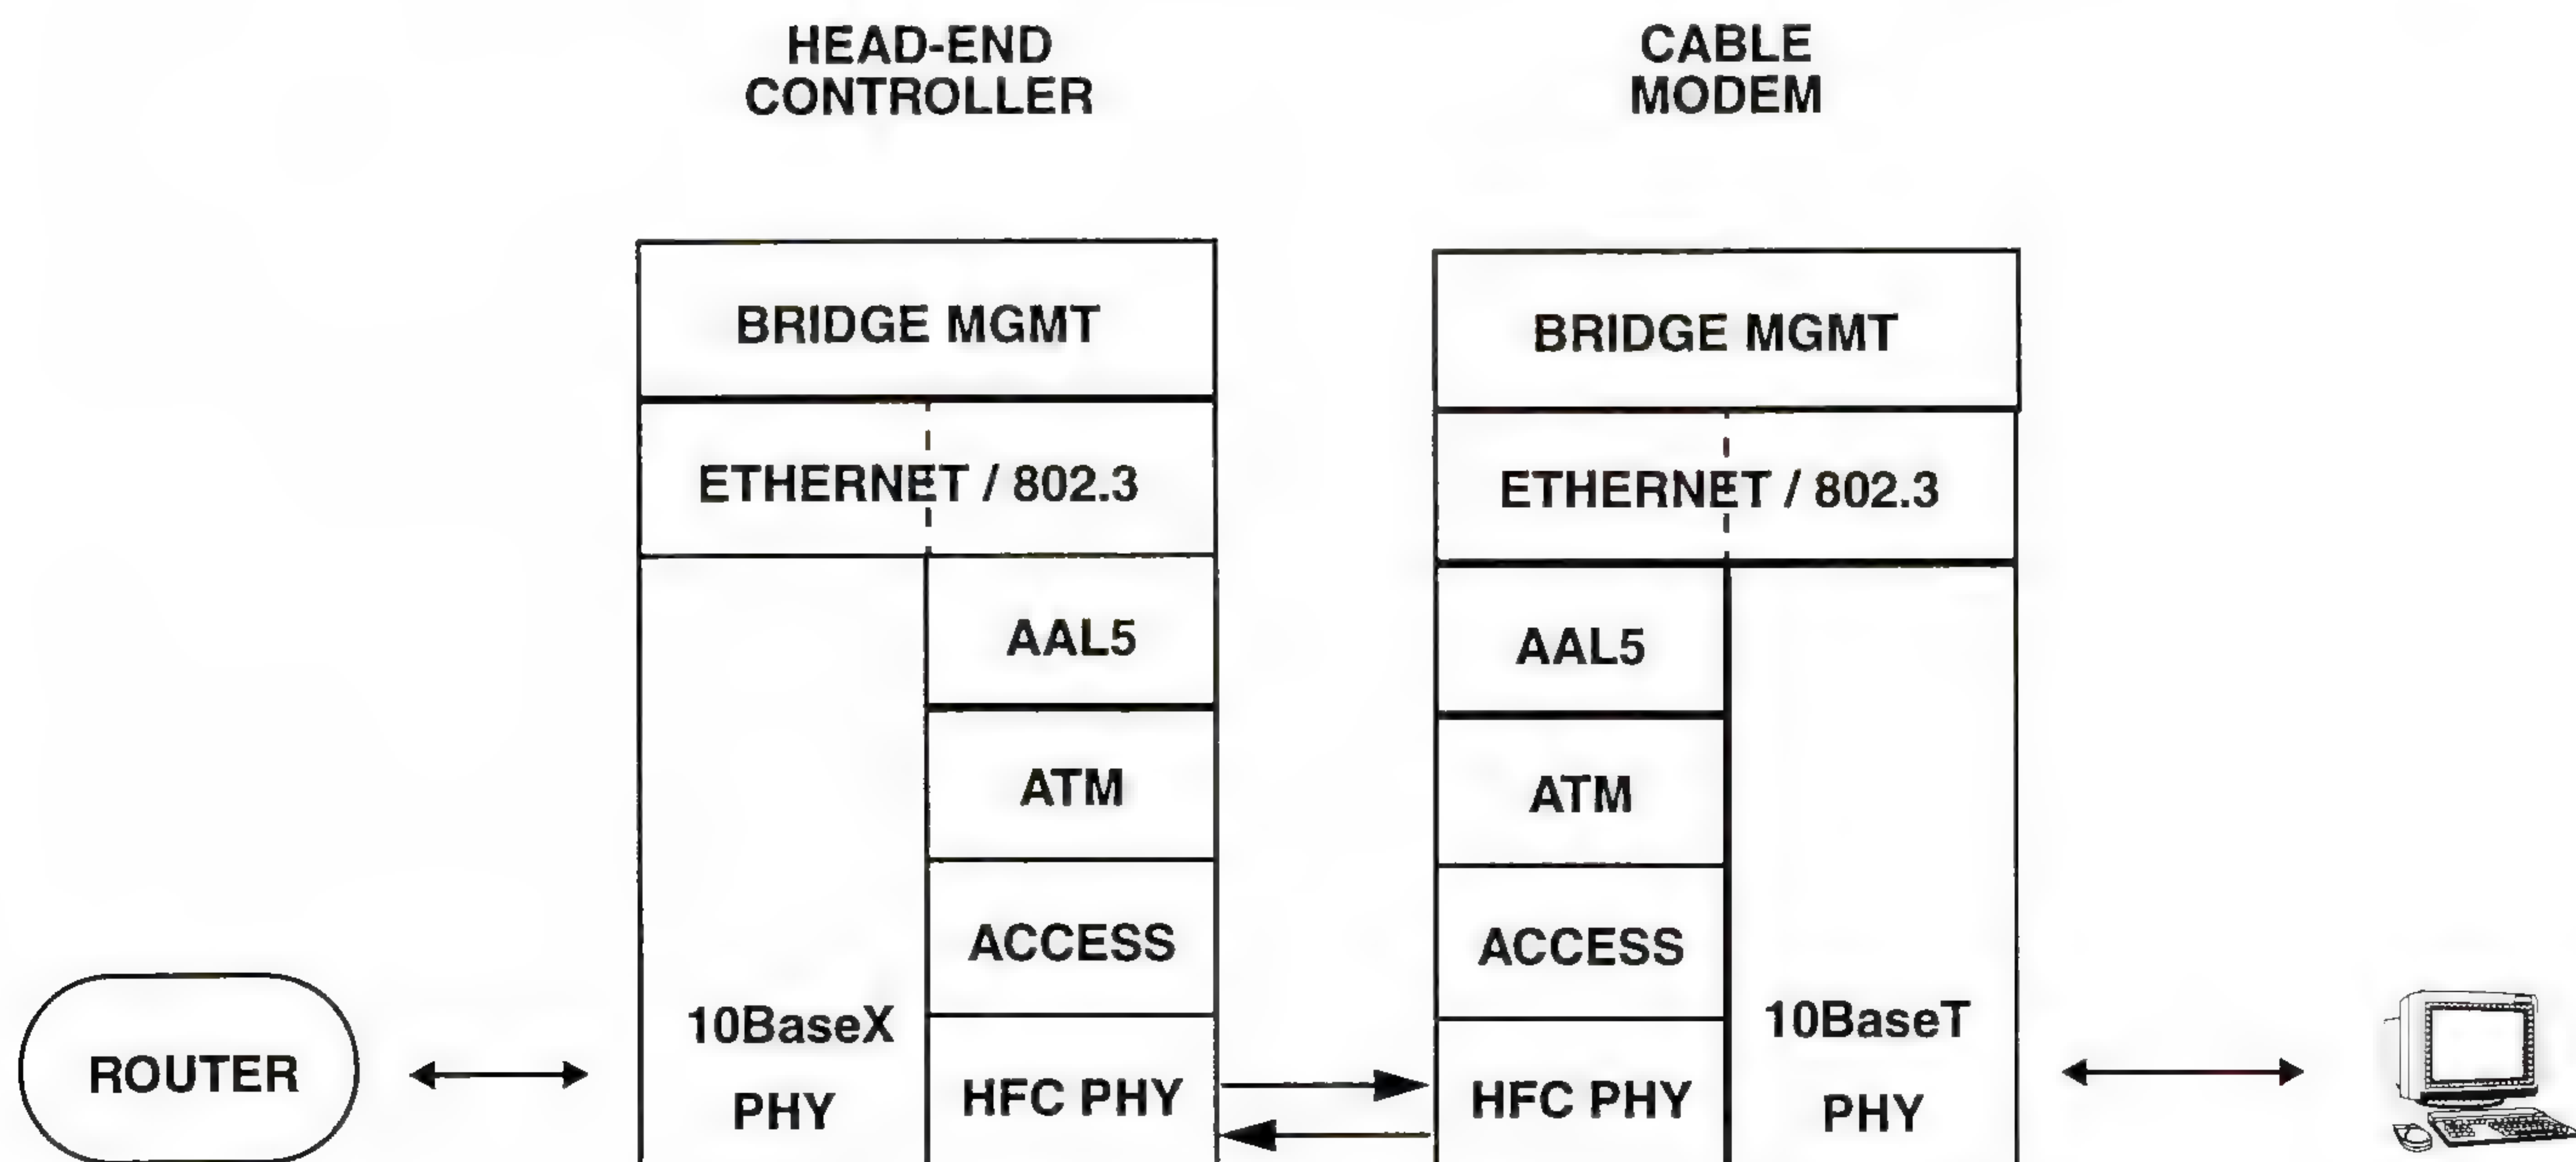


Figure 6: Bridged Ethernet via ATM Example

There are several available methods for layering an Ethernet and 802.3 bridging function [9, 10] over ATM permitting an ATM over HFC system to act as an ATM serial connection between these half-bridge functions. Figure 6 illustrates the protocol stack for this solution. Note that in this configuration, the external user interfaces are Ethernet and not ATM. This feature allows the ATM over HFC system to be viewed as an Ethernet segment. Either model provides an Ethernet-like segment to the cable operator. It is well known how to put together such segments to construct larger internetworked networks. Figure 7 illustrates a protocol stack implementation for non-ATM environments where the head-end controller is an IP router and not a bridge. With either of these solutions, IP datagrams may be routed to the cable modem instead of bridged as in the previous example. Routing creates a different model of how IP addresses are allocated to the home network(s) and where IP protocol stacks terminate. It will be the responsibility of the Internet Service Provider to specify the equipment support requirements for meeting a routed deployment.

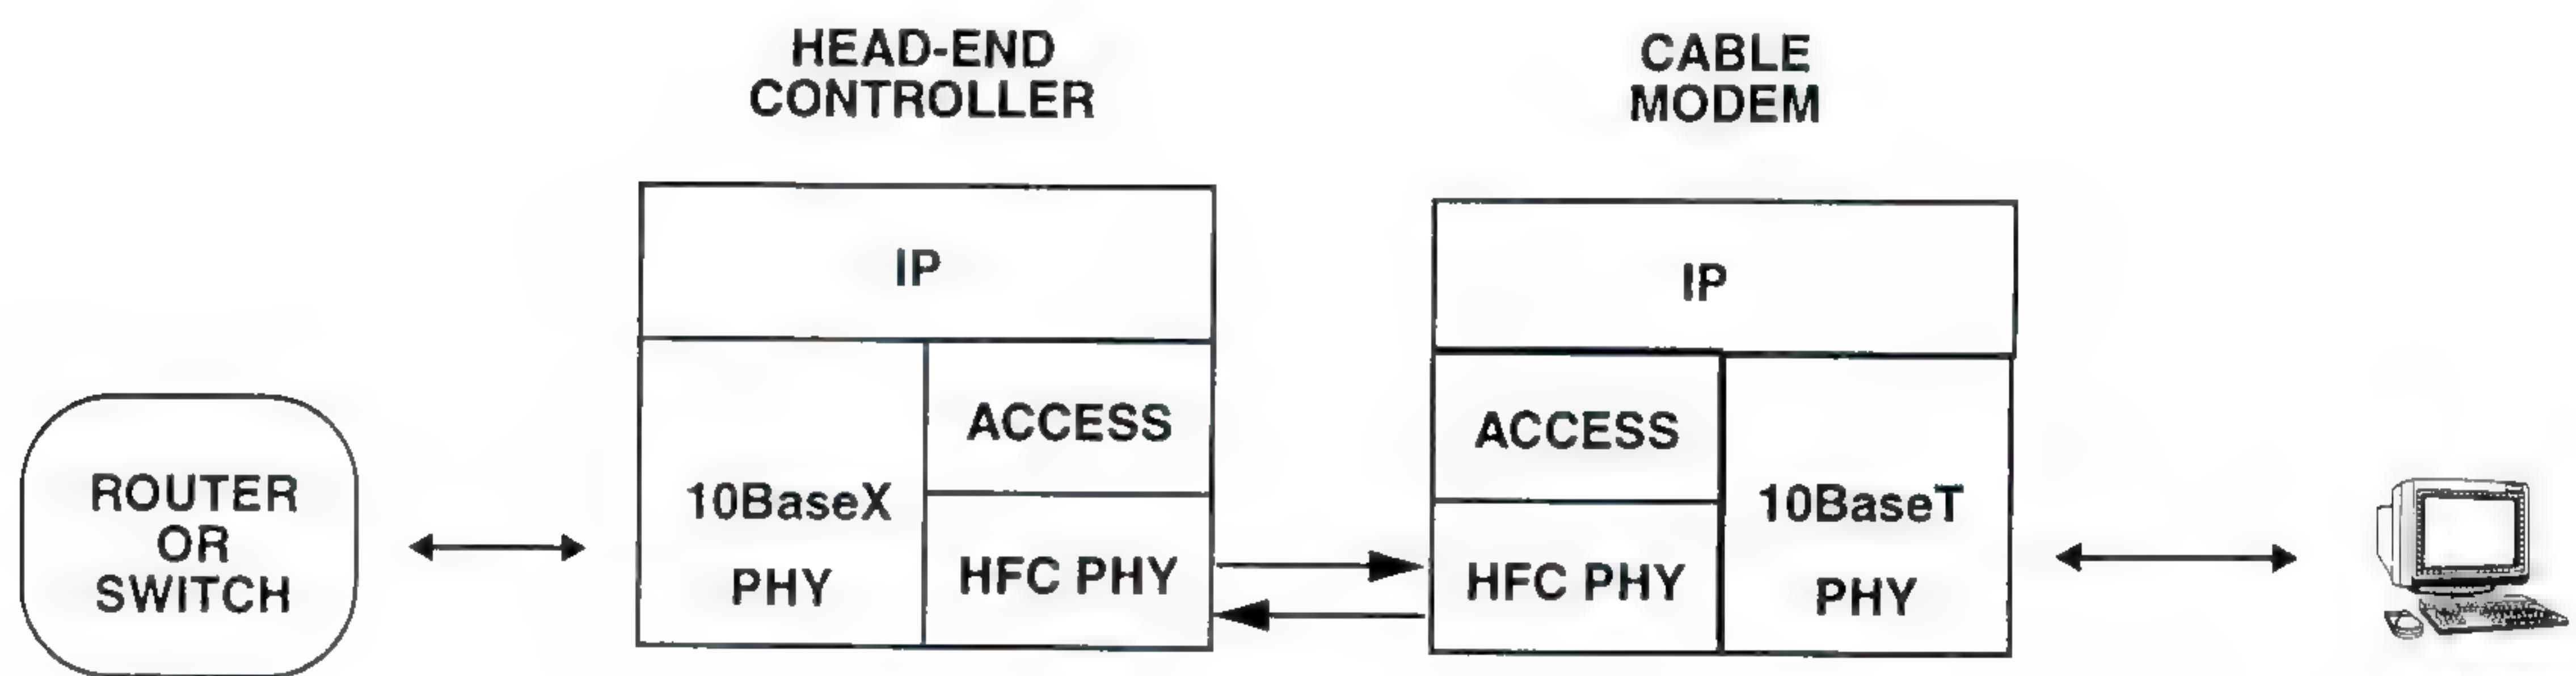


Figure 7: Routed IP Example

Whether a bridging or routing system is employed, the equipment deployment model on the cable network is the same (see Figure 8). In this environment, the cable modems provide a form of service demarcation between the Internet Service Provider's network and each home network.

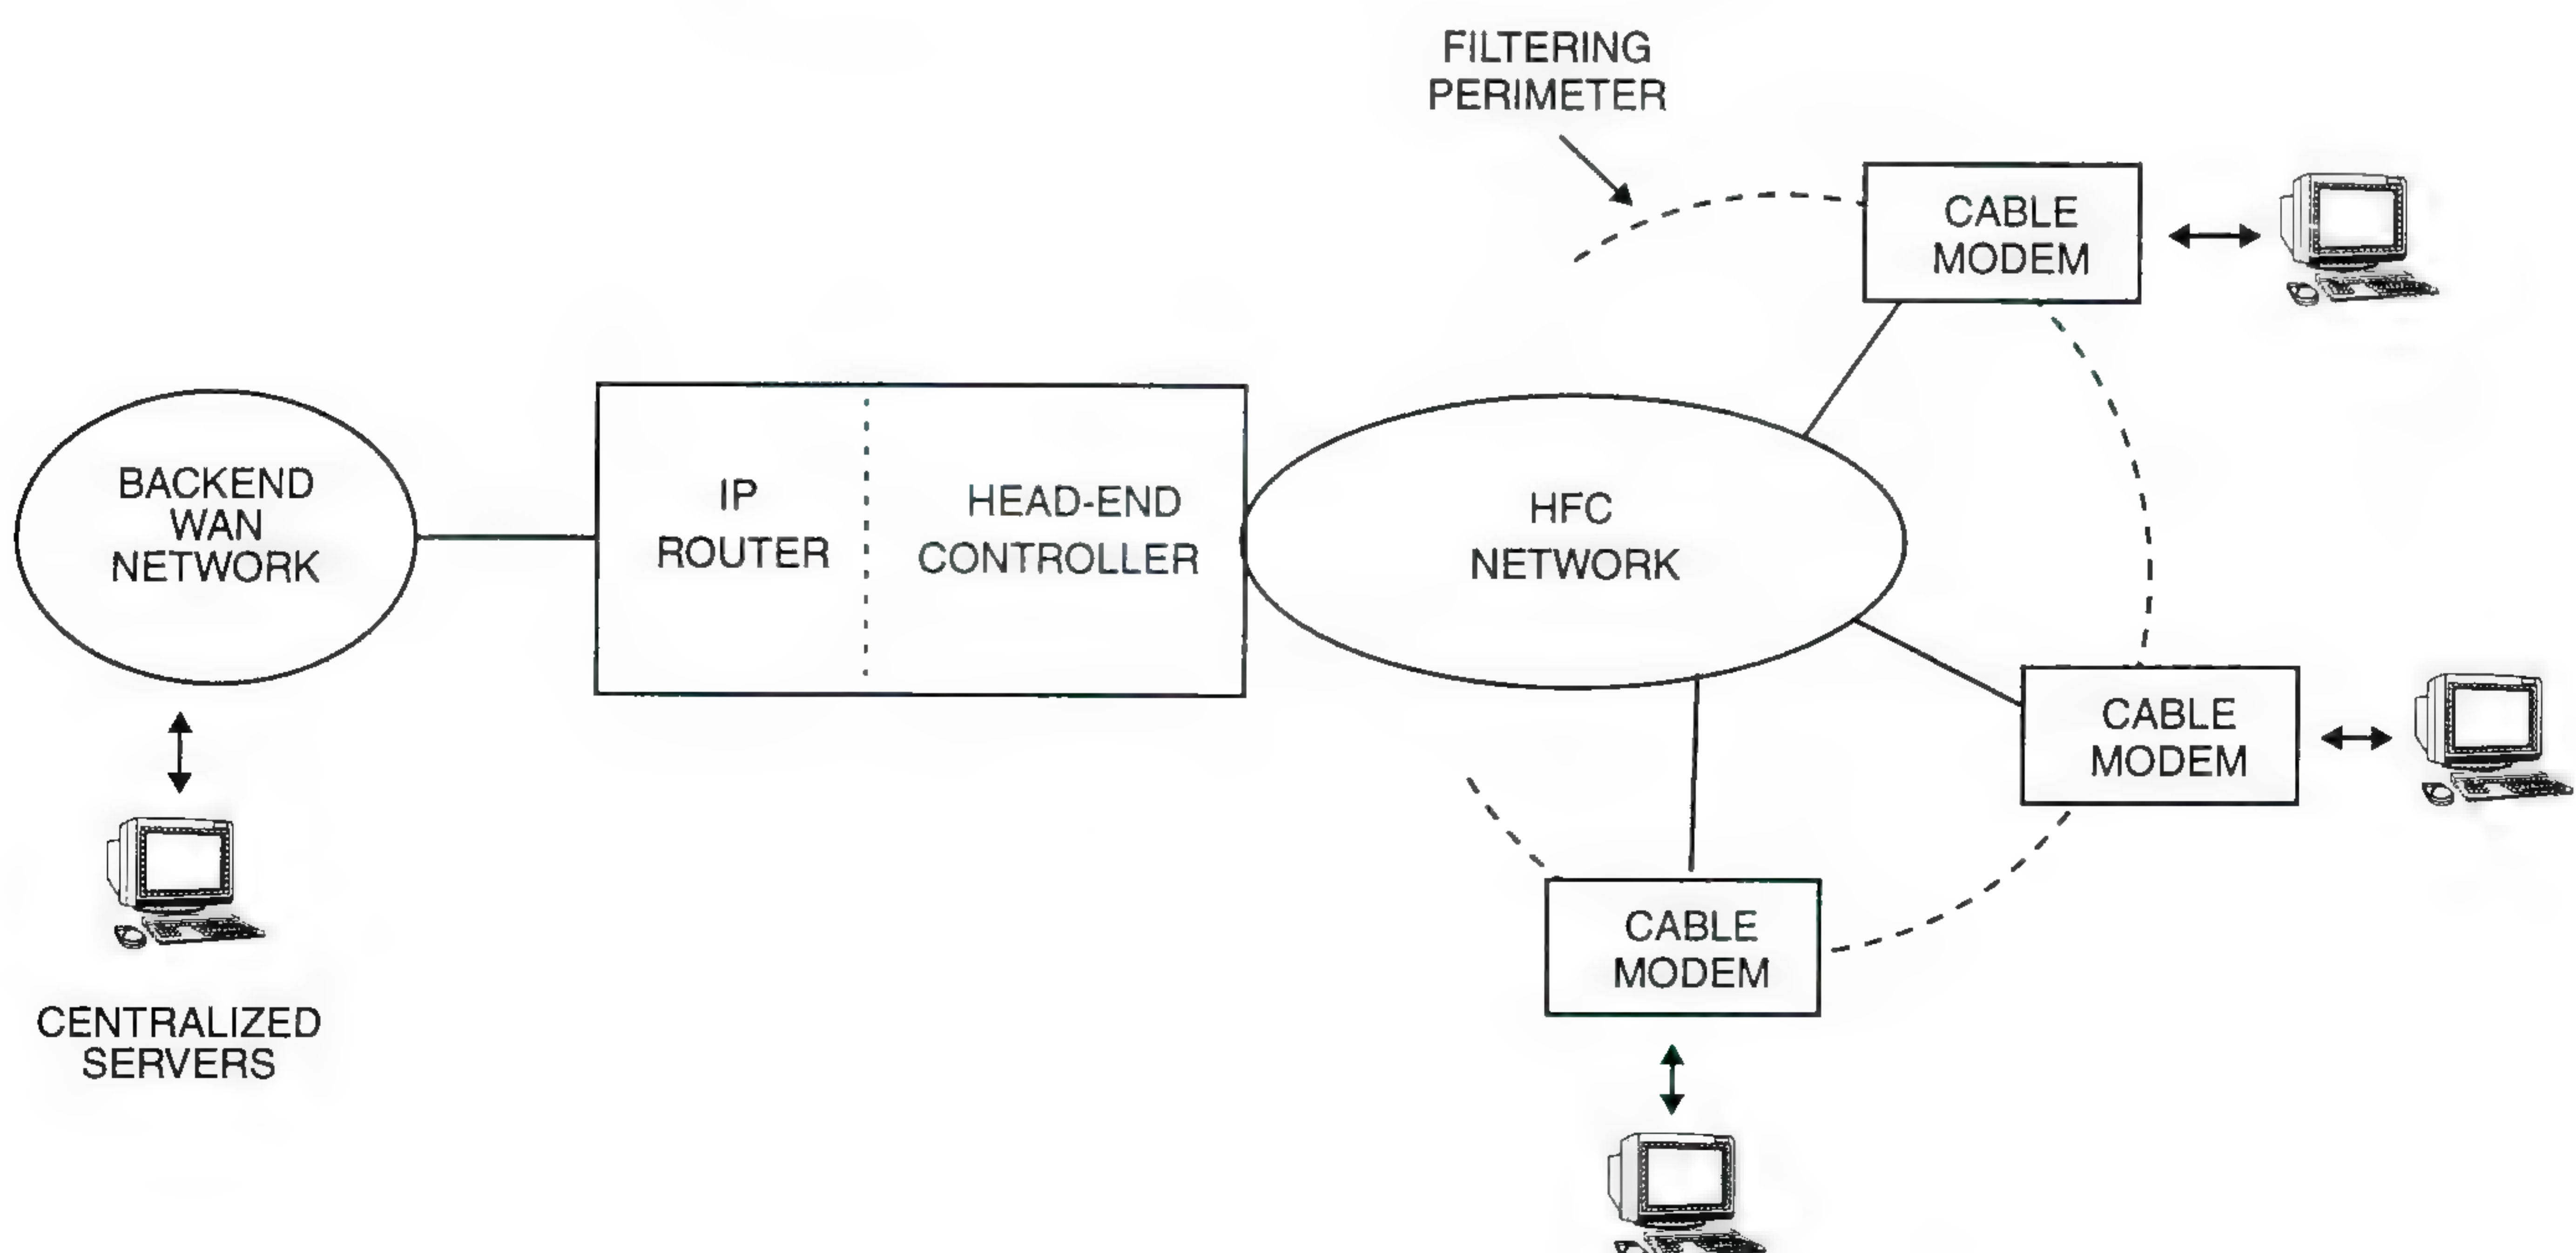


Figure 8: Internet Services via Cable Modem Deployment Model

continued on next page

Residential Broadband Technology (*continued*)

Filtering

To help the Internet Service Provider provide fair access service to its residential customers, the cable modem will require sufficient dynamic functionality for multi-layer protocol filtering and various forms of rate management. The goal of this filter is to create a “defense perimeter” at the first point of entry to the cable network which will protect the upstream channel from being saturated or abused by misbehaving home networks. Some examples of this filtering functionality include, but are not limited to:

- Filtering on Ethertype for permitting only certain protocols to pass upstream: e.g., IP and ARP only
- Filtering on IP source and/or destination address to permit/deny access from the home network.
- IP and Ethernet broadcast rate limiting; i.e., keep any home network broadcast storms confined to the home network
- IP Multicast group address filtering; i.e., explicitly permit participation of the home network in an IP multicast group

It should be noted that these filtering functions are under consideration by a number of cable modem manufacturers. It is expected that the defense perimeter requirements of the Internet service provider and cable operators will evolve over time as these systems are deployed.

This has been a brief overview of the providing IP over cable TV networks. From an engineering and deployment viewpoint making the Internet move over cable modems is deceptively straightforward. There are many issues which are beyond the scope of this article: address allocation methods, back end network design, configuration services, server placement, home customer support services, installation, firewalling, and troubleshooting.

Summary

This article has presented an overview of the work in progress of the ATM Forum’s Residential Broadband Working Group and that of the IEEE P802.14 Cable TV MAC and PHY Protocol standards Working Group. Initial review of these works is positive and indicate that ATM over HFC systems can be constructed using a MAC layer access approach.

Transporting data over Hybrid Fiber-Coax Cable TV networks is technically challenging as the upstream cable plant has numerous hostile noise hurdles which must be overcome by the PHY protocol. These hurdles will require the combined efforts of the standards groups, cable modem manufacturers, and cable operators in order to overcome the problems.

While it is a straightforward matter of engineering to bring all the pieces of the cable network equipment puzzle together, it should be noted that the industry is discovering that data over RF is presenting unforeseen challenges that will take much time and experience in order to sort out. Cable modems are harder to build than anticipated. Cable operators are discovering just how much technical care effort is required to enable and maintain their cable plants for upstream spectrum.

The cable network environment will provide a very usable platform for delivering Internet services to and from the home. A hypothetical example was provided which illustrates a general equipment deployment model.

Actual deployment of Internet to the home will occur in many areas of North America in 1996, and information about these experiences should be available to the community in the near future.

The Internet Service Providers and cable operators should all have this sorted out over the next several years.

Acknowledgements

The author wishes to thank the participants of both the ATM Forum Residential Broadband Working Group and the IEEE P802.14 Cable TV MAC and PHY Working Group for their efforts to date.

This article is dedicated to Foster Bunny, the resident lagomorph at the *ConneXions* home office.

For more information

Information on the IEEE's P802.14 Working Group can be found on the World-Wide Web at:

<http://www.com21.com/pages/ieee-802.14.html>

Information the Internet Engineering Task Force's IP over ATM Working Group can be found at:

<http://www.com21.com/pages/ipatm.html>

The ATM Forum is a closed industrial consortia and non-published work-in-progress documents cannot be distributed publicly to non-members. General information about the ATM Forum may be obtained from the Web at:

<http://www.atmforum.com/>

References

- [1] ATM-Forum, "User Network Interface (UNI) Specification Version 3.0," ISBN 0-13-225863-3, Prentice Hall, December 1993.
- [2] Kim, Albert, "Two-Way Plant Characterization," Technical Session 23, National Cable Television Association Show and Conference, Dallas Texas, 9 May 1995.
- [3] Laubach, Mark, "Classical IP and ARP over ATM," RFC 1577, January 1994.
- [4] IEEE P802.14 Working Group, "Cable-TV Functional Requirements and Evaluation Criteria," work in progress, IEEE802.14/94-002R2, IEEE 802 Committee, February 1995.
- [5] IEEE, "IEEE Standards for Local Area Networks: Logical Link Control," ANSI/IEEE Std 802.2-1985, Fifth Printing, February 1988.
- [6] Abramson, Norman, "Development of the ALOHANET," *IEEE Transactions on Information Theory*, Volume IT-31, pp. 119-123, March 1985.
- [7] Roberts, L., "Extensions of Packet Communication Technology to a Hand Held Personal Terminal," Proceedings of SJCC, pp. 295-269, 1986.
- [8] Baran, Paul, "On Distributed Communication Networks," *IEEE Transactions on Communication Systems*, Volume CS-12, pp. 1-9, March, 1964.
- [9] XEROX, "The Ethernet, A Local Area Network: Data Link Layer and Physical Layer Specification," X3T51/80-50, Xerox Corporation, Stamford, CT, October 1980.

Residential Broadband Technology (*continued*)

- [10] IEEE, "Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications," Standard 802.3-1985 (ISO DIS 8802/3), IEEE, New York, ISBN 0-471-82749-5, 1985.
- [11] Chelehemal, Majid, Prodan, Richard, et al., "Field Evaluation of Reverse-Band Channel Impairments," Society of Cable Telecommunications Engineers, Emerging Technologies Conference, San Francisco, CA, 9-12 January, 1996.
- [12] Laubach, Mark, "ATM for your internet—But When?" *ConneXions*, Volume 7, No. 9, September 1993.
- [13] Atkinson, Ran, "Towards Real ATM Interoperability," *ConneXions*, Volume 7, No. 8, August 1993.
- [14] Laubach, M., "IP Over ATM and the Construction of High-Speed Subnet Backbones," *ConneXions*, Volume 8, No. 7, July 1994.
- [15] Clapp, G. & Zeug, M., "Components of OSI: Asynchronous Transfer Mode," *ConneXions*, Volume 6, No. 4, April 1992.
- [16] Rybczynski, Tony, "The ATM Enterprise Network Switch: Enabler of Network Transformation," *ConneXions*, Volume 9, No. 11, November 1995.
- [17] Heinänen, Juha, "Multiprotocol Encapsulation over ATM," RFC 1483, March 1993.
- [18] Lucien Rhodes, "The Race for More Bandwidth," (Interview with Milo Medin of @Home), *WIRED*, 4.01, January 1996, page 140.

MARK LAUBACH holds a BEE. and M.Sc. from the University of Delaware. He is the Technical Ambassador and Chief Architect at Com21, Inc. in Mountain View California and responsible for the end-to-end systems architecture and ATM over HFC protocol specification of the Com21 product family. Prior to Com21, he was with the Hewlett-Packard Company for 14.5 years. Mark is a member of the IETF and is the chair of the IP-over-ATM Working Group. He is the author of RFC 1577 "Classical IP and ARP Over ATM." Additionally he attends the ATM FORUM and is the appointed liaison from the Forum's Residential Broadband Working Group to the IEEE P802.14 CATV MAC and PHY Protocol Working Group. He is a past member of the CSNET Executive Committee and a past member of the board of trustees for the Corporation for Research and Educational Networking (CREN = CSNET + BITNET). E-mail: laubach@com21.com

Volume 10, 1996 is underway

Included with this issue of *ConneXions* is the Volume 9, 1995 index sheet. If you would like a copy of the 1987-1994 index booklet, just drop us a note. This is also a good time to order your 1996 *ConneXions* binder. The binder costs \$15 including shipping anywhere in the world, order yours today.

Our subscription department has asked me to remind you that when paying for your subscription, you may choose to transfer funds directly to our bank: The Bank of California, San Mateo, CA 94402, ABA = 121000015, account number 027-012797. (This information changed when we switched subscription agencies.) Please include subscriber name and customer number on all wire transfers. For any questions, call 1-800-575-5717 or +1 610-892-1959.

Future NetWorld+Interop Dates and Locations

| | | |
|---------------------|--------------------|-----------------------|
| NetWorld+Interop 96 | Las Vegas, NV | April 1–5, 1996 |
| NetWorld+Interop 96 | Frankfurt, Germany | June 10–14, 1996 |
| NetWorld+Interop 96 | Tokyo, Japan | July 22–26, 1996 |
| NetWorld+Interop 96 | Atlanta, GA | September 16–20, 1996 |
| NetWorld+Interop 96 | Paris, France | October 7–11, 1996 |
| NetWorld+Interop 96 | London, England | Oct. 28–Nov. 1, 1996 |
| NetWorld+Interop 96 | Sydney, Australia | November 25–29, 1996 |
| NetWorld+Interop 97 | Singapore | April 7–11, 1997 |

All dates are subject to change.

More information

Call 1-800-INTEROP or +1-415-578-6900 for more information. Or send e-mail to info@interop.com or fax to +1-415-525-0194. For the latest information about NetWorld+Interop including *N+I Online!* as well as other SOFTBANK produced events, check our home page at <http://www.interop.com>

NetWorld+Interop is produced by SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California 94404–1138, USA.

NETWORLD+INTEROP 96

Write to *ConneXions*!

We'd love to hear your comments, suggestions and questions about anything you read in *ConneXions*. Our editorial address is given below. Use it for letters to the Editor, requests for the index of back issues, questions about particular articles etc.:

ConneXions—The Interoperability Report

303 Vintage Park Drive

Suite 201

Foster City

California 94404–1138

USA

Phone: +1 415-578-6900 or 1-800-INTEROP (Toll-free in the USA)

Fax: +1 415-525-0194

E-mail: connexions@interop.com

URL: <http://www.interop.com>

Subscription information

For questions about your subscription please call our customer service hotline: 1-800-575-5717 or +1 610-892-1959 outside the USA. This is the number for our new subscription agency, Seybold Publications. Their fax number is +1 610-565-1858. The mailing address for subscription payments is: P.O. Box 976, Media, PA 19063–0976.

This publication is distributed on an "as is" basis, without warranty. Neither the publisher nor any contributor shall have any liability to any person or entity with respect to any liability, loss, or damage caused or alleged to be caused, directly or indirectly, by the information contained in *ConneXions—The Interoperability Report*®

CONNEXIONS

303 Vintage Park Drive
Suite 201
Foster City, CA 94404-1138
Phone: 415-578-6900
FAX: 415-525-0194

ADDRESS CORRECTION
REQUESTED

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD

Dr. Vinton G. Cerf
Senior Vice President, MCI Telecommunications
President, The Internet Society (1992 – 1995)

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute



Printed on recycled paper

Subscribe to CONNEXIONS

U.S./Canada ☐ \$195. for 12 issues/year

All other countries ☐ \$245. for 12 issues/year

Name _____ Title _____

Company _____ E-mail _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

Fax () _____

☐ Check enclosed (in U.S. dollars made payable to **CONNEXIONS**).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card# _____ Exp. Date _____

Signature _____

Please return this application with payment to:

Back issues available upon request \$15./each
Volume discounts available upon request

CONNEXIONS

303 Vintage Park Drive, Suite 201
Foster City, CA 94404-1138
415-578-6900 FAX: 415-525-0194
connexions@interop.com

CONNEXIONS